The American Journal of Interdisciplinary Innovations and Research
ISSN 2642-7478 | Open Access

Type Original Research
PAGE NO. 91-100
DOI 10.37547/tajiir/Volume07Issue07-08

The Modern CPG Data Stack: Building End-To-End Analytics on Azure, Snowflake, And Dbt

Supriya Gandhari
Independent Researcher , USA

OPEN ACCESS

SUBMITED 11 June 2025
ACCEPTED 28 June 2025
PUBLISHED 11 July 2025
VOLUME Vol.07 Issue 07 2025

CITATION
Supriya Gandhari. (2025). The Modern CPG Data Stack: Building End-To-End Analytics on Azure, Snowflake, And Dbt. The American Journal of Interdisciplinary Innovations and Research, 7(07), 91–100. https://doi.org/10.37547/tajiir/Volume07Issue07-08

Abstract:

Consumer Packaged Goods (CPG) firms encounter distinct obstacles in handling the extensive data produced from sales, supply chains, customer interactions, and market patterns. To tackle these obstacles, organizations are progressively implementing a contemporary data stack. This document examines how using Microsoft Azure, Snowflake, and DBT technologies can create a comprehensive analytics stack that improves scalability, data transformation, and decision-making processes. The paper outlines the technical framework, best practices, and real-world examples that showcase the implementation of this modern CPG data stack. By utilizing cloud-based solutions, businesses can enhance operational efficiency, automate data processes, and acquire more profound insights. We will investigate performance metrics, security risks, cost-efficiency methods, and forthcoming trends that influence the advancement of CPG analytics.

Keywords: Modern Data Stack, DBT, Snowflake, Azure Data Factory, Data Engineering, Cloud Data Platforms, ETL Pipelines, Data Transformation, Data Orchestration, Data Integration, Real Time Analytics, Cloud Native Solutions

1. INTRODUCTION

In recent years, the swift digital transformation of companies has resulted in a significant increase in data volume and complexity. Conventional data management systems have found it challenging to adapt to these changes, leading to the emergence of a more adaptable

and scalable framework— often termed the Modern Data Stack (MDS). The MDS encompasses a collection of cloud-native tools that work together to facilitate effective data ingestion, storage, transformation, analysis, and governance [1]. This shift is altering the way organizations extract value from data and is especially impactful in data-centric industries such as e-commerce, finance, healthcare, and technology. A graphical representation of modern data stack has been shown in Fig. 1.

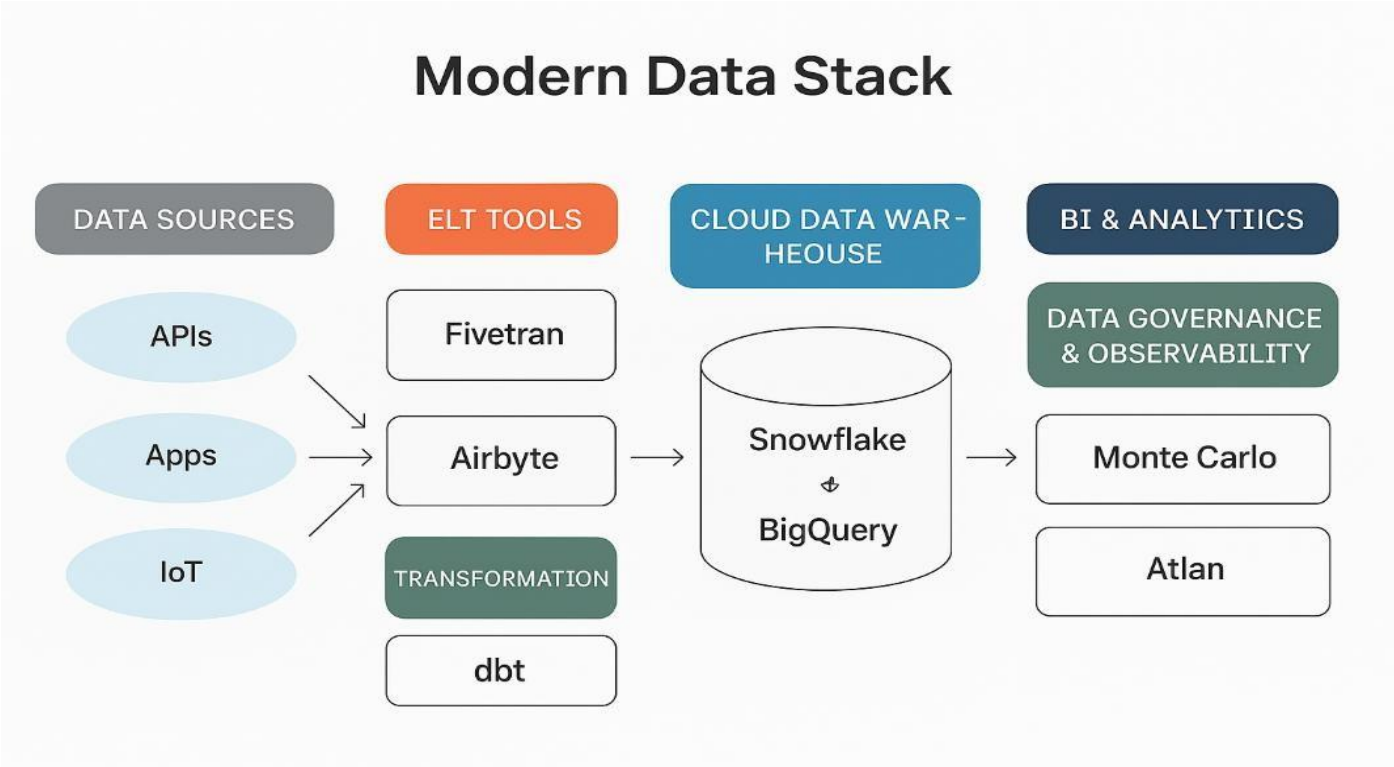

Fig. 1. Modern Data Stack Architecture

The Consumer-Packaged Goods (CPG) sector produces vast amounts of data from various sources such as retail sales, online shopping platforms, logistics systems, and customer interaction points. Traditionally, businesses depended on outdated systems that were isolated, sluggish, and unable to manage contemporary data requirements [2]. Nevertheless, due to innovations in cloud technology, data storage, and transformation tools, companies are now transitioning to a completely unified analytics framework.

Even with the growing utilization of the modern data stack, there is a notable absence of thorough academic research into its architecture, advantages, and obstacles. Most available literature centers on specific components (such as cloud data warehouses or ELT tools) rather than providing a comprehensive perspective on the entire stack [1]. Consequently, the practical integration and strategic deployment of MDS components are poorly understood in many businesses. This paper seeks to address this gap by providing an in-depth analysis of the modern data stack and comparing it to traditional data systems regarding scalability, performance, and cost-effectiveness. This research is particularly pertinent at a time when organizations are reassessing their data strategies due to increasing data privacy regulations, the necessity for real-time analytics, and the rise of third-party SaaS applications.

The main goals of this paper are:

To define and articulate the components of the modern data stack. To investigate the primary factors driving the transition from traditional data systems to contemporary architecture.

To assess the practical consequences and effects on organizations resulting from the implementation of a modern data stack.

To emphasize ongoing challenges and future avenues for research and practice.

This study offers a structured synthesis of existing tools and practices within modern data infrastructures, backed by both industry examples and academic insights.

Microsoft Azure offers a cloud-based infrastructure that allows for efficient storage, computation, and analytics. Snowflake delivers a robust, cloud-native data warehouse capable of handling large- scale data processing with flexibility. DBT (data build tool) acts as a framework for orchestration that empowers data analysts and engineers to efficiently transform and model data within the warehouse [2]. Collectively, these technologies establish the backbone of a contemporary CPG data stack that supports real-time analytics, predictive modeling, and automation. This research paper provides an in-depth exploration of the implementation and enhancement of this data stack. It addresses architecture design, data transformation strategies, security considerations, and practical case studies. By following best practices in data engineering, CPG companies can enhance their analytics stack for scalability, cost-effectiveness, and performance.

Grasping the modern data stack is vital for businesses looking to remain competitive in the digital landscape. Its importance lies in facilitating agile decision-making, reducing costs through scalable cloud infrastructure, and promoting data accessibility across different business units. By integrating perspectives from both industry and academia, this paper aims to assist stakeholders— from data engineers to executives—in making well-informed choices regarding investments in data infrastructure.

2. Why a Modern Data Stack?

The consumer-packaged goods (CPG) sector encounters distinct data-related obstacles, including high volumes of transactions, intricate supply chains, and omnichannel sales. Outdated data systems, typically reliant on on-premises SQL databases or legacy ETL processes, lack the flexibility needed to meet contemporary analytical demands. Adopting a modern cloud-based data architecture offers several advantages like:

- **Scalability:** The capability to manage extensive datasets effectively.
- **Real-time Processing:** Swift insights enabled by the computational power of the cloud.
- **Cost Efficiency:** Pay-as-you-go pricing structures help lower infrastructure expenses.
- **Interoperability:** Smooth integration across various tools and data sources.

By utilizing Azure for infrastructure, Snowflake for data warehousing, and DBT for data transformation, CPG firms can establish a robust analytics system that is adaptable, scalable, and economically viable.

The Modern Data Stack (MDS) is a flexible ecosystem of cloud-based tools and technologies that collaborate with:

- Collect
- Store
- Clean
- Transform
- Analyze
- Manage data from various dispersed sources.

This framework updates inflexible legacy data systems, providing agility, adaptability, and a cohesive pipeline from raw data to usable insights. By making data accessible to more users, the MDS grants both technical and non-technical individuals the ability for self-service analytics, while also facilitating strong data governance measures. Consequently, organizations can uphold confidence in their data, adhere to regulatory standards, and guarantee that insights are both prompt and precise. A key component of this stack is the cloud data warehouse or data lake, which consolidates vast amounts of structured and unstructured data while allowing for real-time processing and scalability [4]. The stack also features elements for data transformation (such as DBT), workflow orchestration (Such as Dagster), and data visualization (like Power BI, Looker, and Tableau)—all integrated to enhance AI/ML applications and business intelligence efforts (Fig. 2).

At its essence, the contemporary data stack offers:

- ⚡ Agility through real-time data processing and swift implementation
- 🔁 Automation throughout the complete data lifecycle
- 🔐 Governance and quality backed by metadata tracking and observability

. 📊 Insights at scale, supporting both human decision-making and AI-driven actions.

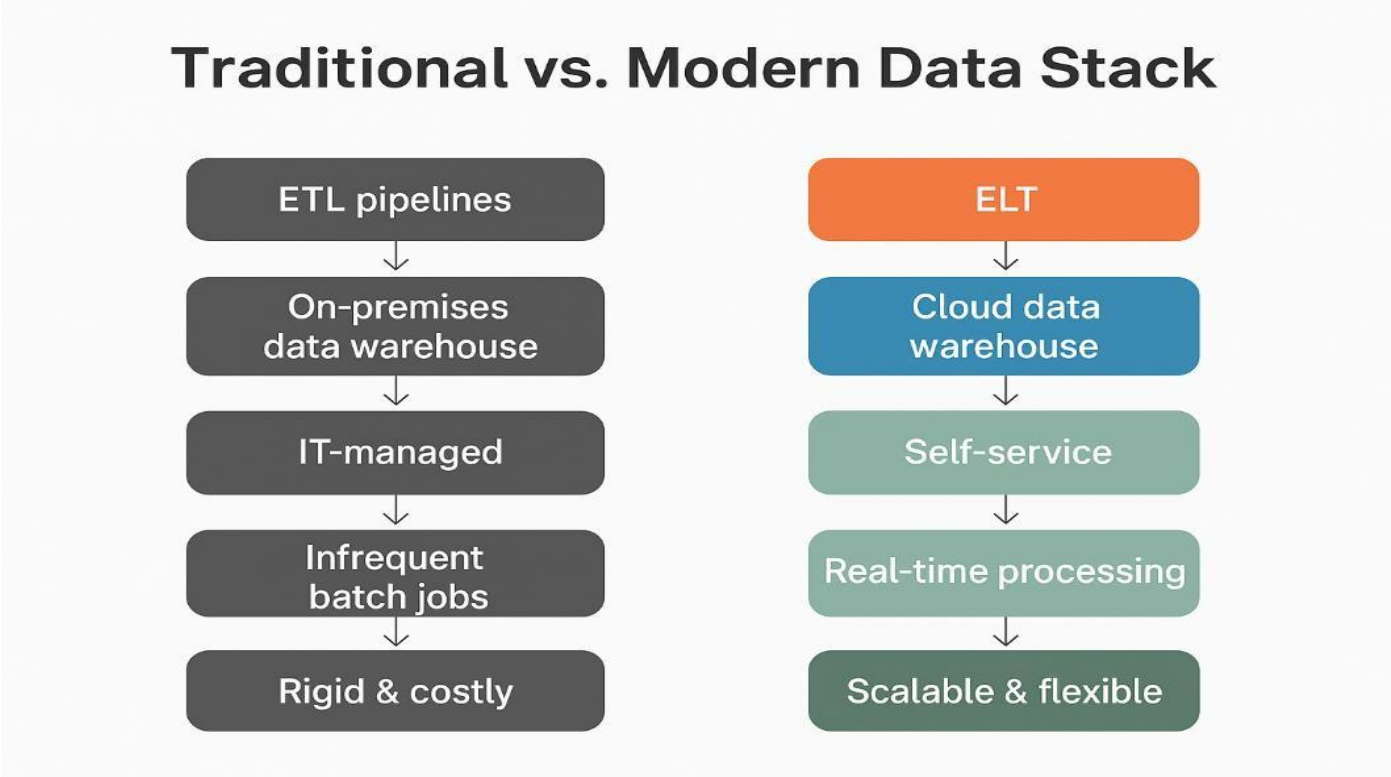

Fig. 2. Traditional vs Modern Data Stack

This framework promotes a robust data culture, stimulates innovation, and serves as the foundation of digital transformation—ensuring organizations that can swiftly and consistently derive value from data. Whether managing large datasets, streaming events, or conventional records, the modern data stack equips you with the infrastructure, tools, and adaptability needed to succeed in today's data-oriented landscape [5].

3. Azure: The Foundation for Cloud Infrastructure

The Consumer-Packaged Goods (CPG) sector is facing heightened demands to adapt quickly, leverage data, and foster innovation in response to swift changes in consumer preferences, supply chain fluctuations, and competitive pressures. To tackle these challenges, CPG companies are transitioning from outdated, on-premises systems to adopting the modern data stack — a versatile, cloud-based framework that focuses on scalability, automation, immediate access, and integration with artificial intelligence [4].

Azure Data Factory (ADF) offers powerful features for seamless integration with Snowflake, enabling effective data movement and transformation. Key benefits include:

Incremental Data Loading: ADF enables incremental data extraction by utilizing Snowflake's change tracking capabilities. This process ensures that only the data altered since the last run of the pipeline is processed, thereby improving performance and resource efficiency.

Change Data Capture (CDC): ADF can make use of Snowflake's inherent change tracking to handle only the delta data since the last pipeline run. This is accomplished by establishing a Stream object on the source table and employing the CHANGES clause to retrieve newly added or modified data.

Query Pushdown Optimization: By setting up ADF pipelines to reduce unnecessary intermediate transformations, you can better harness Snowflake's processing capabilities, which boosts performance for substantial datasets.

Secure Authentication: ADF has the ability to connect to Snowflake via key-pair authentication, which bolsters security by removing the requirement to store passwords.

Microsoft Azure acts as a fundamental platform in this transition, offering the cloud infrastructure necessary to support comprehensive data workflows on a scale. Azure provides storage options like Data Lake Storage Gen2, computing services such as Azure Synapse

Analytics, and integration functionalities with enterprise systems like ERP, CRM, and outside data sources [2]. Its global reach, top-tier security, and compliance features make it a suitable choice for multinational CPG enterprises.

Central to the architecture of the modern data stack within Azure is Snowflake, a cloud-native data warehouse solution recognized for its unique separation of storage and computing capabilities. Snowflake operates effortlessly on Azure, allowing companies to consolidate their extensive data—from production and logistics to customer transactions and marketing analytics—into a single, easily accessible repository. Its compatibility with semi-structured data formats (e.g., JSON, Parquet), features like time travel, and secure data sharing capabilities facilitate quicker and more collaborative data analysis [7].

To guarantee that data is not merely stored but is instead converted into a clean, analytics-ready format, DBT (data build tool) is commonly incorporated within the stack. DBT empowers data teams to carry out SQL-based transformations directly within Snowflake, utilizing a modular, version-controlled, and testable development framework. This approach significantly enhances data reliability and clarity, as each transformation step is documented and reproducible. DBT also allows teams to implement software engineering best practices within data, including automated testing, CI/CD pipelines, and lineage tracking, all of which are essential for governance and long- term sustainability.

Together, Azure, Snowflake, and DBT equip CPG organizations with a fully integrated ecosystem for ingesting, transforming, and analyzing data (Fig. 3). Azure Data Factory or other ingestion tools can channel raw data from various sources into Snowflake. DBT executes the transformation logic directly within the warehouse, and the resulting datasets are utilized in downstream analytics platforms, frequently integrated with Azure services like Power BI or Azure Machine Learning for reporting and predictive analytics.

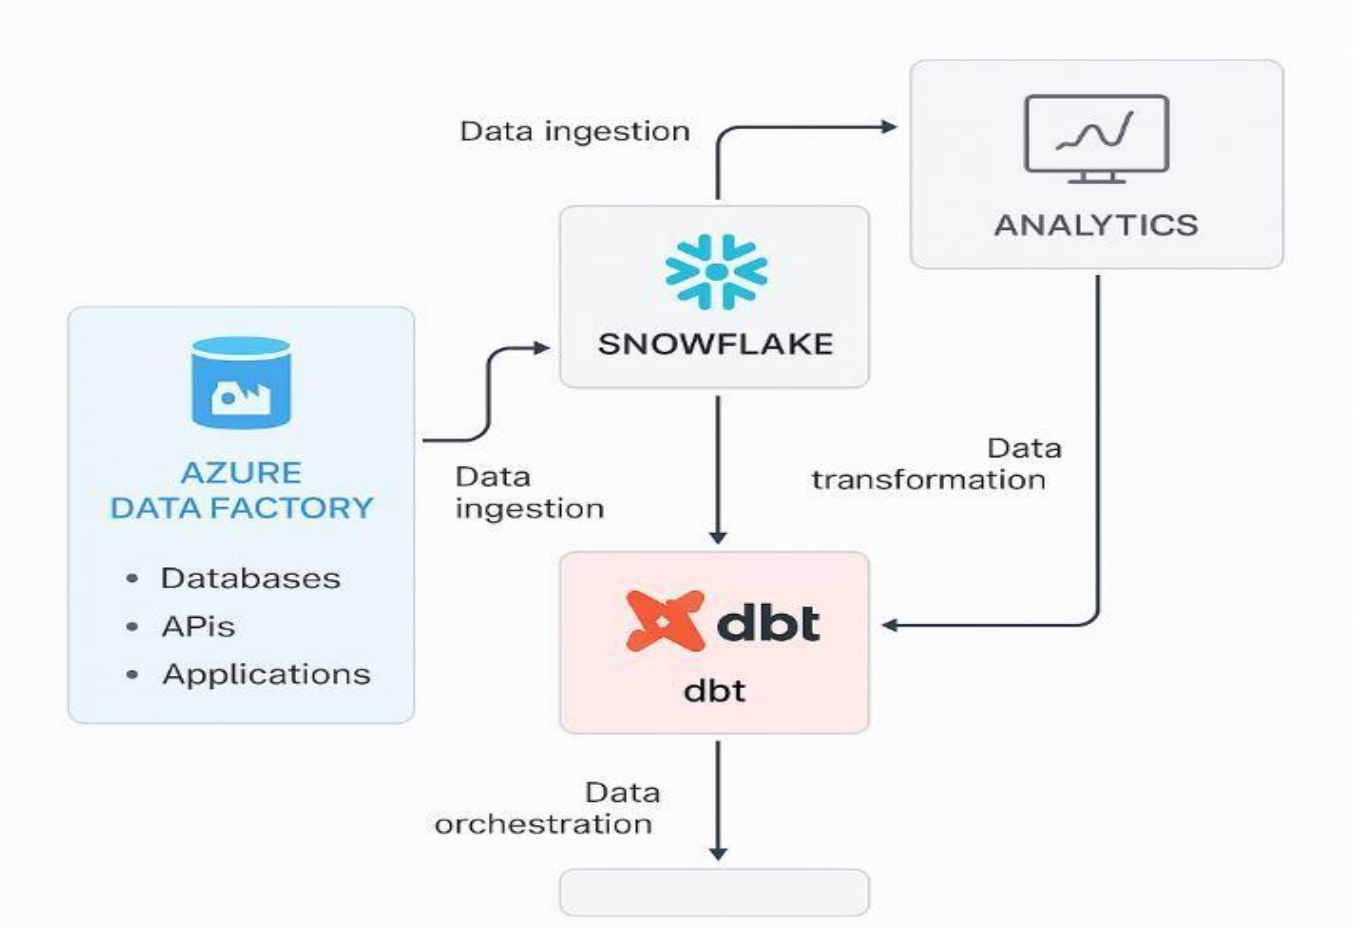

Fig. 3. Data Flow Pipeline

This modern data stack framework benefits the CPG sector by providing near real-time visibility into the supply chain, improving forecasting and inventory control, accelerating product development timelines, and enhancing personalized marketing efforts. Moreover, it enables both technical and non-technical users to access and investigate data, promoting a culture centered on data-driven decision-making. The scalability offered by Azure, the performance and versatility of Snowflake, and the structured transformation abilities of dbt collectively create a system that addresses the intricate data requirements of CPG organizations while minimizing operational costs and boosting responsiveness.

In summary, the modern data stack powered by Azure, Snowflake, and DBT enables CPG firms to integrate their data, optimize operations, support AI-driven innovations, and produce business insights rapidly. It signifies a pivotal evolution from batch-oriented, siloed analytics to a more real- time, flexible, and intelligent approach to data management.

4. Snowflake: The Scalable Cloud Data Warehouse

Snowflake is a prominent cloud-based data warehouse platform designed to offer quick, adaptable, and scalable data management for contemporary organizations. As sectors increasingly embrace digital transformation tactics, Snowflake has become a fundamental component of the modern data ecosystem—facilitating real-time analytics, AI integration, and global data collaboration within intricate data environments.

Unlike traditional on-premises or legacy data warehouses, Snowflake was specifically developed for the cloud. It functions on major cloud platforms such as Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP), giving customers the option to select or implement a multi-cloud approach [6]. Its distinct architecture separates storage and compute, enabling users to scale each one independently according to their requirements. This separation guarantees that large-scale analytical tasks can run simultaneously without affecting performance, a significant benefit in data-heavy industries like Consumer-Packaged Goods (CPG).

Snowflake natively supports both structured and semi-structured data, including formats such as JSON, Avro, Parquet, and XML. This capability enables businesses to consolidate various data types—customer transactions, supply chain telemetry, market research, and IoT signals—into a single data warehouse, where it can be queried effectively using SQL. Its features like time travel and zero-copy cloning are also streamline version control and experimentation, which are vital in collaborative data settings.

A key feature of Snowflake is its capacity to enable secure and seamless data sharing across different business units, partners, and geographical locations. Through its Data Sharing functionalities, organizations can distribute live, read-only datasets with external parties without the need for data migration or duplication. This is particularly beneficial for coordinating supply chains, third-party reporting, and compliance-driven data visibility in sectors like CPG, where collaborative efforts across organizations are prevalent [5].

For data transformation and analytical processes, Snowflake integrates directly with contemporary tools such as DBT (data build tool), allowing transformations to occur within the data warehouse using modular SQL. This close integration improves transparency, supports version control, and simplifies deployment workflows. Analysts and engineers can create, test, and document data models in a collaborative setting, ensuring that business logic remains consistent and repeatable across various teams.

Snowflake also addresses modern analytics needs such as real-time reporting, predictive modeling, and AI/ML workloads. With its support for external functions, stored procedures, and integrations with machine learning platforms, Snowflake serves as a central hub for both conventional business intelligence and advanced data science. Additionally, features like Snowpark allow developers to create custom data processing logic in programming languages such as Python, Java, and Scala, enhancing the platform's versatility and reach [3].

Governance, compliance, and security are integral to Snowflake's architecture. Features such as fine-grained access control, role-based permissions, end-to-end encryption, and built-in auditing ensure that data is safeguarded and accessible only to authorized users. These elements are particularly essential in regulated industries or for global companies that need to adhere

to data protection regulations like GDPR and CCPA.

In summary, Snowflake revolutionizes how organizations store, manage, and utilize data. By delivering elastic scalability, high performance, and seamless integration with the broader modern data stack, it fosters quicker decision-making, improved operational efficiency, and lays the foundation for future-ready data initiatives. Whether managing extensive data volumes or intricate analytics tasks, Snowflake's cloud-native architecture empowers teams to leverage data without being hindered by infrastructure constraints [8].

5. DBT: The Transformation Layer for Analytics

The consumer-packaged goods (CPG) sector functions in an environment that is both intensely competitive and constantly changing, where timely insights and quick decision-making are essential. Data is gathered from a variety of sources—such as retail point-of-sale systems, supply chain platforms, digital marketing tools, and ERP systems—and must be consolidated, transformed, and presented promptly and reliably. This is where DBT (data build tool) becomes crucial in the contemporary data ecosystem [5].

DBT is a transformation tool that operates within the cloud data warehouse itself—most prominently on platforms like Snowflake, which is often implemented on Microsoft Azure infrastructure. It enables data teams to create modular SQL transformations and model their data in layers, akin to how software engineering teams assemble codebases. These transformations are treated as code, allowing teams to version-control, test, document, and deploy transformation logic consistently and effectively.

dbt allows data analysts and engineers to more effectively transform data within the warehouse. Key practices include:

Modular and Reusable Code: By organizing dbt projects into modular components, teams can enhance code reusability and maintainability across various projects.

Incremental Models: dbt enables the use of incremental models that process only newly added or modified records, significantly cutting down on processing time and reducing computational expenses.

Built-in Testing and Documentation: dbt facilitates the implementation of tests to ensure data integrity and the creation of documentation for data models, improving transparency and governance.

CI/CD Integration: Connecting dbt with Continuous Integration and Deployment pipelines guarantees consistent and reliable execution of data transformations across different environments.

In the framework of the modern data ecosystem, DBT is positioned between the raw data ingestion layer (managed by tools like Azure Data Factory or ingestion services linked to Snowflake) and the analytics or machine learning layers (such as Power BI, Tableau, Azure ML, etc.). Its role is to clean, join, filter, and reshape data to prepare it for analysis. DBT accomplishes this through a method called ELT (Extract, Load, Transform), where raw data is first centralized in Snowflake and then transformed within the warehouse.

The reason DBT is especially advantageous for CPG is its capacity to manage complexity in a scalable and maintainable manner. CPG data environments tend to be extensive, featuring thousands of SKUs, various regions, multiple currencies, and intertwined customer and product hierarchies. DBT enables teams to encapsulate business logic into reusable, auditable SQL models. For example, one model could outline "net sales by brand," while another might standardize "customer segments" or "retail delivery windows." This structured, modular method minimizes errors and accelerates iteration cycles [4].

DBT also facilitates data validation through its built-in testing features. CPG organizations can formulate assertions to verify data quality conditions, such as the presence of null values in key fields, the uniqueness of product identifiers, or referential integrity between transactional and master data. These tests are performed alongside transformations, helping to establish trust in the data before it is delivered to business stakeholders or machine learning workflows.

When integrated within an Azure ecosystem, DBT integrates particularly well with CI/CD workflows utilizing Azure DevOps or GitHub Actions. This enables CPG data teams to deploy adjustments to transformation logic in a controlled, automated, and repeatable manner. Together with Snowflake's performance and scalability—alongside Azure's

orchestration, security, and compliance features, DBT forms a strong foundation for enterprise-level data transformation.

Furthermore, DBT promotes a data culture characterized by transparency and thorough documentation. Every model developed in DBT can contain comprehensive descriptions, connections to other models, and lineage tracking, which DBT automatically generates into interactive documentation. This feature is particularly vital in large CPG organizations where collaboration across various functions is necessary, and the business logic needs to be shared and comprehended among marketing, supply chain, finance, and data science teams.

As CPG firms increasingly integrate AI and real-time analytics, the transformation layer becomes even more essential. DBT enables the creation of feature stores for machine learning, the preparation of datasets for forecasting models, and the development of dashboards with consistent logic—all derived from a single, authoritative source of truth.

Essentially, DBT acts as the transformation engine that links raw, often chaotic data to refined, reliable datasets that facilitate reporting, predictive analytics, and strategic decision-making within the CPG sector. Its integration with Snowflake on the Azure cloud empowers CPG businesses to function at scale, retain agility, and enforce data governance—all crucial elements in a dynamic, data-intensive industry [7].

In the realm of the CPG industry, where information is sourced from various platforms—from retail point-of-sale systems and marketing tools to ERP and supply chain databases—creating a responsive, scalable, and intelligent analytics environment is essential for business success. The combination of Azure, Snowflake, and dbt into a unified, modern data stack architecture allows CPG companies to accomplish this objective by simplifying data ingestion, transformation, analysis, and governance.

The process starts with Microsoft Azure, which provides essential cloud infrastructure. Azure offers a secure and scalable platform for ingesting and storing large amounts of data, whether structured or unstructured. CPG companies frequently utilize Azure Data Factory (ADF) to manage data movement from various source systems such as SAP (ERP), Salesforce (CRM), manufacturing logs, retail distributor feeds, and e-commerce platforms. ADF facilitates either scheduled or real-time ingestion of this data into a centralized storage solution, typically Azure Data Lake Storage Gen2 [9].

Once the raw data arrives in the data lake, it is imported into Snowflake, which serves as the cloud- native data warehouse and main repository for analytics. Snowflake, running on Azure, offers scalability in both compute and storage, which is especially beneficial in CPG settings that experience high variability in data and intricate reporting requirements. Since Snowflake distinguishes between compute and storage, CPG organizations have the flexibility to adjust data workloads separately allowing them to execute complex queries and dashboards without interrupting active data loads or batch processing tasks.

With the raw data now available in Snowflake, the transformation layer is managed by DBT (data build tool). DBT functions directly on Snowflake's compute engine, enabling data teams to create modular SQL models that convert raw tables into datasets ready for analytics. This transformation process includes data cleansing, deduplication, standardization of dimensions such as product hierarchies or customer segments, and intricate metric definitions like net revenue, margin contribution, and demand forecasts [8]. DBT allows CPG analysts and engineers to define these transformations as code under version control, applying software engineering practices like modularity, testing, and documentation.

The DBT layer is where raw operational data is transformed into standardized business logic. For instance, a sales data model might consolidate daily transactions across various geographies and retailers, standardize currency conversions, and connect product SKUs to marketing initiatives. Each transformation step is meticulously tracked and documented, guaranteeing transparency and reproducibility. DBT's integrated testing features validate key assumptions—ensuring referential integrity, checking for nulls in vital fields, or identifying unexpected data spikes. This is crucial in the CPG sector where decisions often rely on rigorously controlled metrics [10].

The well-structured and reliable transformed data in Snowflake becomes the foundation for diverse analytics

and decision-making processes. Business intelligence platforms like Power BI or Tableau establish direct connections to Snowflake to deliver interactive dashboards and real-time reporting for teams across marketing, finance, logistics, and executive leadership [2]. Furthermore, Azure Machine Learning or external ML platforms can extract curated datasets from Snowflake to train and deploy predictive models for applications like sales forecasting, price elasticity modeling, or customer churn prediction.

Azure enhances architecture by providing governance, security, and operational automation. Azure DevOps facilitates CI/CD pipelines for DBT projects, allowing code-based transformations to be automatically deployed with peer reviews and version control. Azure Key Vault safeguards credentials and secrets, ensuring secure access to Snowflake and other system components. Azure Monitor and Log Analytics deliver centralized oversight of data workflows and system performance [4].

The result of this comprehensive architecture is a fully modernized data ecosystem. Data transitions from ingestion to insight through an automated and transparent pipeline. Business logic

is ingrained in maintainable code. Scalability is inherent in cloud-native infrastructure. Most importantly, business users across the CPG organization gain quicker, more reliable access to the data they require—empowering decisions that encompass promotional planning, inventory management, and new product launches.

This integrated modern data stack not only enhances the efficiency and precision of analytics but also allows CPG companies to weave data into the foundation of their operational strategy. It facilitates movement from reactive to proactive decision-making and from isolated data management to a collaborative, enterprise-wide data intelligence framework.

6. CONCLUSION

The Modern Data Stack (MDS) has transformed data management by providing modular, scalable, and cloud-based solutions that improve data processing and analysis. Its implementation has allowed businesses to extract actionable insights more effectively, promoting data-driven decision- making.

Nonetheless, as the MDS progresses, various challenges and avenues for further exploration have surfaced:

a. **Integration Complexity and Maintenance Burden:** The emergence of numerous specialized tools within the MDS has resulted in significant integration difficulties. Organizations frequently find themselves dealing with many tools, each presenting unique interfaces and data format requirements. This complexity not only extends the initial setup duration but also complicates day- to-day management, with each integration point representing a possible failure point in the infrastructure. Additionally, upkeep of these tools demands considerable resources, diverting attention from primary business goals.

b. **Alignment Between Data Engineering and Data Science**: Although the MDS has mainly concentrated on data engineering, there is an increasing necessity to connect data engineering with data science. At present, tools aimed at data science, such as feature engineering and model management, are not yet as developed within the MDS ecosystem. Future investigations should look into creating unified platforms that promote better collaboration between these areas.

c. **Data Governance and Quality Control:** Ensuring data integrity and compliance is becoming more difficult as organizations gather and process immense volumes of data. Discrepant data formats and fragmented tools can create data silos and impede cooperation. Research into automated governance solutions and standardized protocols could offer remedies for these challenges.

d. **Integration of Predictive Analytics and Machine Learning:** The emergence of predictive analytics and machine learning within the MDS brings both opportunities and challenges. While these technologies can provide deeper insights, they also introduce intricacies concerning model management and data pipeline integration. Future research should aim to develop frameworks that streamline the integration of advanced analytics into the MDS.

In conclusion, while Modern Data Stack has greatly improved data management practices, continuous

research and innovation are crucial to address its existing limitations and to fully harness its capabilities in various organizational environments.

REFERENCES

1. Kimball, R., & Ross, M. (2013). The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling (3rd ed.). Wiley.
2. Giebler, C., Grunert, A., Saake, G., & Szamél, M. (2019). Big Data Processing with Apache Spark in the Cloud. Journal of Big Data, 6(1), 1–22. https://doi.org/10.1186/s40537-019-0200-7
3. Snowflake Inc. (2022). Modern Data Architecture for the Cloud. Retrieved from https://resources.snowflake.com
4. Raj, P., & Raman, A. (2018). Cloud-native data warehouse architecture: A study of Snowflake computing. In Proceedings of the 5th International Conference on Data Engineering and Communication Technology (pp. 503–511). Springer.
5. Fishtown Analytics. (2020). What is dbt (data build tool)? Retrieved from https://docs.getdbt.com/docs/introduction
6. Microsoft Azure. (2023). What is Azure Data Factory?. Retrieved from https://learn.microsoft.com/en-us/azure/data-factory/introduction
7. Gartner. (2022). Magic Quadrant for Cloud Database Management Systems. Retrieved from https://www.gartner.com
8. Redpoint Ventures. (2021). The Rise of the Modern Data Stack. Retrieved from https://tomtunguz.com/modern-data-stack
9. dbt Labs. (2023). How dbt fits into the Modern Data Stack. Retrieved from https://www.getdbt.com/blog/
10. Deloitte. (2021). Data modernization: How to build a modern data architecture. Retrieved from https://www2.deloitte.com