

# Advanced Analytics with Snowflake and Power BI

Srinivasa Rao Karanam

Srinivasarao.karanam@gmail.com

New Jersey, USA

**Abstract:** The demand for advanced analytics and near real-time insights has intensified across industries. Organizations are increasingly turning to modern data platforms and powerful business intelligence tools to streamline their analytical processes and drive decision-making. Snowflake has cemented its position as a leading cloud data platform that enables elastic compute scaling, secure data sharing, and powerful governance features. Simultaneously, Microsoft's Power BI continues to evolve as one of the most widely adopted business intelligence and analytics solutions. The synergy between Snowflake and Power BI is stronger than ever, combining Snowflake's scalable compute architecture and data governance capabilities with the robust visualization and modeling features of Power BI. By exploring these topics, one can understand why a Snowflake–Power BI integration can provide a powerful and flexible solution for modern data analytics. This article assumes a baseline knowledge of both Snowflake and Power BI, although key concepts are introduced for clarity.

**Keywords:** Snowflake, Power BI, Star Schema, Dynamic Tables, Data Governance, Snowflake Marketplace, Row-Level Security, Data Masking, Near Real-Time Analytics, Cloud Data Warehousing

## I. IMPORTANCE OF ADVANCED ANALYTICS

Companies handle data of ever-increasing volume, velocity, and variety, ranging from transactional data spanning hundreds of millions of records to unstructured data such as text logs, sensor metrics, and more. At the same time, data engineering teams are tasked with delivering analytics pipelines that can adapt quickly to shifting market conditions. Decision-makers require insights that can reflect a near real-time view of business performance, while also demanding granular governance rules to protect sensitive data.

Snowflake, as a multi-cloud data platform, has responded to these challenges by enhancing features like Dynamic Tables for easy incremental transformations, row-access and masking policies for robust security, and the Snowflake Marketplace for fast access to third-party datasets. Power BI has also evolved substantially. It now includes improved DirectQuery performance, refined dataflows, and more advanced aggregation features. These improvements address the real-time refresh needs of enterprise users who wish to minimize the complexities of caching massive datasets in memory.

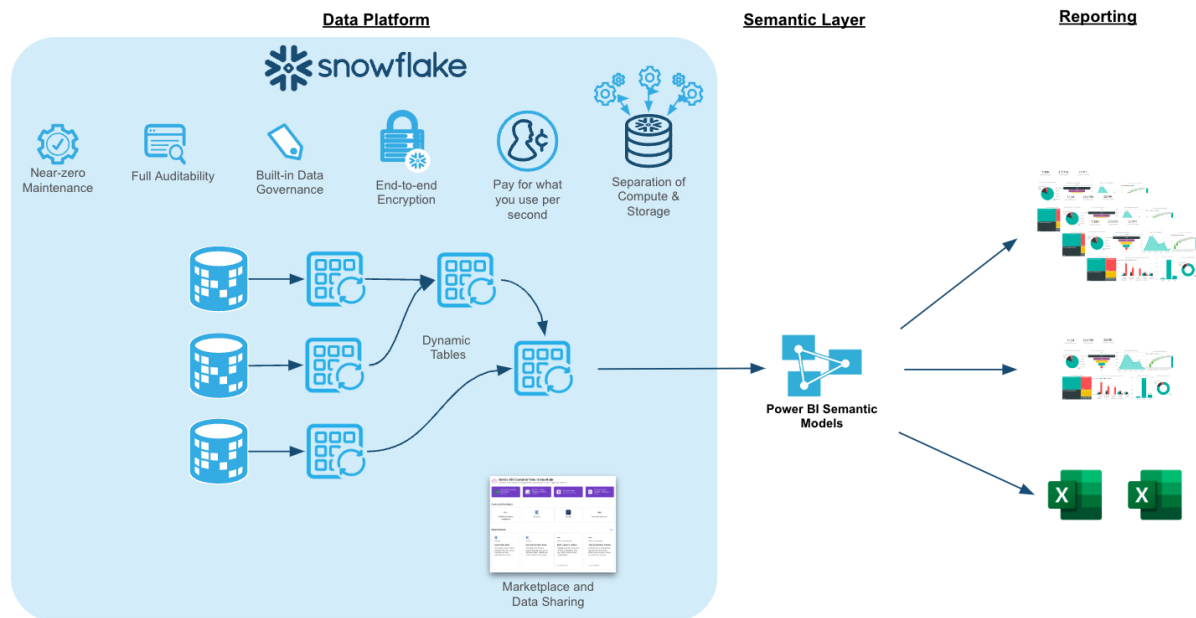


Figure 1: Illustration of Snowflake's data platform integrating with Power BI for semantic modeling and reporting, enabling dynamic tables, data sharing, and visualization for business insights.

In many organizations, data pipelines and analytics workflows still rely on multiple data repositories, siloed operational systems, or legacy on-premises databases. Translating these systems into one cohesive data model is a perpetual challenge, and performance can become a bottleneck when dealing with extremely large fact tables or complex transformations. The Snowflake–Power BI integration, however, has proven to be a strong approach for bridging data storage, transformation, and analytics. Snowflake’s platform provides a secure, governed, and highly performant environment for storing and processing data, while Power BI enables intuitive, self-service analytics that can be adapted to different user personas, from data scientists to business executives.

With this context established, the rest of the article walks step by step through a realistic example scenario, simulating an organization that sells food truck meals under the brand “Tasty Bytes.” We will see how raw point-of-sale (POS) data can be profiled and transformed within Snowflake, enriched with third-party location data, modeled in a star schema using Dynamic Tables, secured with column- and row-level data governance features, and finally connected to Power BI for business intelligence. By observing these steps, data professionals can replicate or adapt the concepts to their own projects.

## II. REVIEWING AND PROFILING THE DATASET IN SNOWSIGHT

In a typical analytics project, one of the earliest steps is to understand the nature of the raw data by conducting a thorough data profiling exercise. When building upon an existing dataset or a previously completed quickstart (for instance, the fictional “Introduction to Tasty Bytes” quickstart), it is important to examine row counts, data distributions, potential data quality issues, and relationships among tables.

Snowsight, Snowflake’s integrated user interface, offers a user-friendly experience for running SQL queries, visualizing results, checking query performance, and collaborating with other data team members. Snowsight has evolved to incorporate more advanced profiling features, interactive charting options, and better integration into other Snowflake functionalities such as Dynamic Tables and Marketplace.

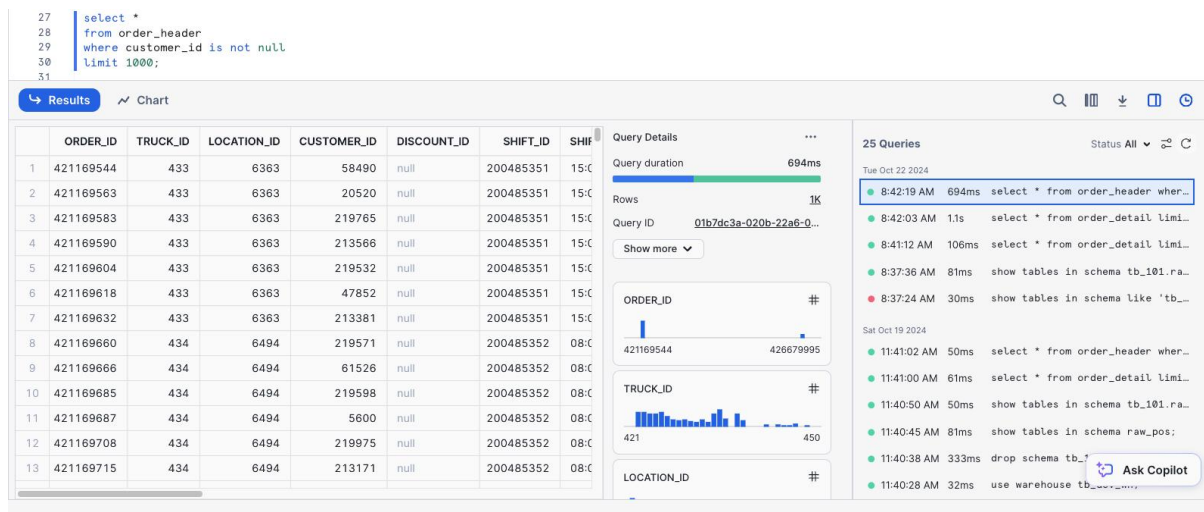


Figure 2: Reviewing a dataset in Snowflake, displaying query execution results, metadata, and performance insights for efficient data analysis and validation.

A typical dataset for Tasty Bytes might include hundreds of millions of rows in order detail tables and tens of millions of rows in order header tables. Given the data's large volume, simply performing naive queries without attention to cardinalities, table sizes, or data quality can lead to inefficient analytics pipelines downstream.

In a research or real-world setting, a user might open Snowsight, navigate to "Projects" and then "Worksheets," and create a specific folder to organize all the SQL scripts for the project. This approach ensures traceability of transformations and queries, critical for both collaboration and auditability. After creating a new folder, the user can create a new SQL worksheet specifically for data profiling tasks.

In an example scenario, the user might set the worksheet context using something like:

```
use role sysadmin;
use database tb_101;
use schema raw_pos;
use warehouse tb_dev_wh;
```

This ensures that the queries run with the correct permissions, target the right database and schema, and use the correct warehouse. The user might next issue queries to see row counts:

```
show tables in schema tb_101.raw_pos;
```

Such a query reveals the cardinalities of each table. Suppose the order detail table has over 670 million records, the order header table holds over 248 million, and the franchise table is significantly smaller. Sampling data in Snowsight can provide quick insights into column distributions, as shown when a user executes a simple **SELECT** statement and reviews the distribution charts in the results pane.

Data profiling also typically includes checks for duplicates. A user may create a common table expression (CTE) to group by a unique ID and identify duplicates where the count is greater than 1, then join back to the main table to locate the specific offending rows. This reveals that certain franchise records might be duplicated or inconsistent, thus informing future data cleaning tasks.

At the close of the data profiling phase, data engineers have a good understanding of table sizes, duplication issues, data distributions, and possible dimension attributes for upcoming transformations. Profiling is more than just a housekeeping step; it forms the foundation for designing efficient schemas and deciding how best to approach transformations.

### III. CREATING NEW SCHEMAS, ROLES, AND WAREHOUSES FOR BUSINESS INTELLIGENCE TEAMS

Once data profiling is complete, the next step is to set up the necessary Snowflake infrastructure for analytics. In a large or even moderately sized organization, multiple teams often need to connect to Snowflake. The scenario described here simulates a global BI analyst team, and thus we need new roles and a warehouse specifically for Power BI usage.

Snowflake's role-based access control architecture is crucial for ensuring data security and preventing unauthorized access. Rather than giving all users elevated privileges, a best practice is to create discrete roles representing different responsibilities. For instance, in this Tasty Bytes scenario, one might create roles such as `tb_bi_analyst_global`, `tb_bi_analyst_na`, `tb_bi_analyst_emea`, and `tb_bi_analyst_apac`. These roles can then be granted to a user `tb_bi_analyst` who will act as a stand-in for distinct regional analysts in testing.

A new schema named `tb_101.powerbi` might also be created to store the curated and transformed tables that will power the BI reports in Power BI. A typical snippet for creating this schema is as follows:

```
create or replace schema tb_101.powerbi;  
use schema tb_101.powerbi;
```

In this example, the data engineer might also create a new warehouse called `tb_powerbi_wh`. This warehouse can be sized appropriately for the queries that Power BI will send in DirectQuery mode, a decision that typically accounts for concurrency and the complexity of the expected queries. The script might look like:

```
create or replace warehouse tb_powerbi_wh  
warehouse_size = 'MEDIUM'  
max_cluster_count = 1  
min_cluster_count = 1  
auto_suspend = 300  
initially_suspended = true  
comment = 'Warehouse used for the TB Power BI DQ semantic model';
```

Setting up dedicated warehouses for specific workloads remains a best practice in Snowflake. This approach can isolate performance between different workloads and simplify cost chargeback. For example, if the data

engineering transformations run on one warehouse and the Power BI usage occurs on another, each team can be better held accountable for its usage and performance.

The creation of new roles in the Snowflake ecosystem also highlights the importance of a layered security model. The commands to create roles, assign them to a user, and grant them privileges on both the schema and the warehouse are often repeated tasks in real-world pipelines:

```
create or replace role tb_bi_analyst_global;
create or replace role tb_bi_analyst_emea;
create or replace role tb_bi_analyst_na;
create or replace role tb_bi_analyst_apac;

grant role tb_bi_analyst_global to user tb_bi_analyst;
...
grant usage on warehouse tb_powerbi_wh to role tb_bi_analyst_global;
...
```

At the end of these steps, the environment is ready for subsequent transformations. The data engineer has established separate staging, raw, and curated schemas, along with specialized roles and a dedicated warehouse for analytics.

#### IV. ENRICHING DATA WITH THIRD-PARTY DATASETS FROM THE SNOWFLAKE MARKETPLACE

To drive deeper analytics, organizations increasingly utilize external data sources. Snowflake Marketplace simplifies this process by allowing Snowflake users to search for and “subscribe” to third-party datasets. In many cases, these datasets are shared in such a way that no additional copy is required. The consumer can join data directly from the shared database to their tables, thus eliminating many of the complex steps previously associated with external data ingestion.

In this Tasty Bytes scenario, location-based data from SafeGraph exemplifies how quickly data can be augmented. The user navigates to the Snowflake Marketplace interface via the left-hand panel in Snowsight and searches for a listing such as “SafeGraph: Frostbyte.” After reading the overview, the user clicks “Get,” granting specific roles access to the shared database. Moments later, the user finds that a new shared database is available in their Snowflake environment. The user then executes cross-database joins from the newly added `safegraph_frostbyte.public` schema to the main Tasty Bytes location table. For instance, if both the Tasty Bytes location table and the SafeGraph data contain a `placekey` column, the following type of query might be used:

```
select
  l.location_id,
  l.city as location_city,
  sg.street_address as location_street_address,
  sg.postal_code as location_postal_code,
  sg.latitude as location_latitude,
  sg.longitude as location_longitude
from tb_101.raw_pos.location l
```

```
left join safegraph_frostbyte.public.frostbyte_tb_safegraph_s sg  
on sg.placekey = l.placekey;
```

By the end of this step, data teams can see the new location attributes or address details integrated with their existing POS data. Some users may choose to replicate this shared data into their own schema for transformations or further enhancements. Others may opt for ongoing cross-database queries if real-time updates from the provider are important. Either way, the frictionless combination of new data in Snowflake underscores the platform's utility in building advanced analytics solutions.

## V. TRANSFORMING DATA INTO A STAR SCHEMA WITH DYNAMIC TABLES

Analytical projects often require data transformations before the data is exposed to visualization tools. While some transformations can be done within a BI tool's data layer, best practices usually dictate that all significant transformations, particularly those concerning large volumes of data, occur within the data platform itself. This approach provides centralized governance, ensures consistent logic across multiple use cases, and leverages the scale of cloud platforms like Snowflake.

Power BI, like many BI tools, performs best when data is arranged in a star schema. A star schema typically consists of a central fact table containing numeric measures (and keys to dimensions) and multiple dimension tables providing descriptive attributes of the measures. For instance, in the Tasty Bytes scenario, the fact table might contain millions of sales order lines, each with foreign keys referencing dimensions such as Customer, Truck, Location, and Date. Having each dimension in a separate table leads to simpler relationships, improved query performance, and a more intuitive self-service experience for analysts.

In the 1990s, Ralph Kimball popularized dimensional modeling, and those principles are still very much relevant. Power BI, for instance, can leverage star schemas to produce more efficient DAX queries. One of the most tangible benefits is that dimension tables usually contain fewer rows, which makes them quick to query for slicers and filters. Meanwhile, the fact table often has fewer columns but significantly more rows, storing transactional data at varying grains of detail.



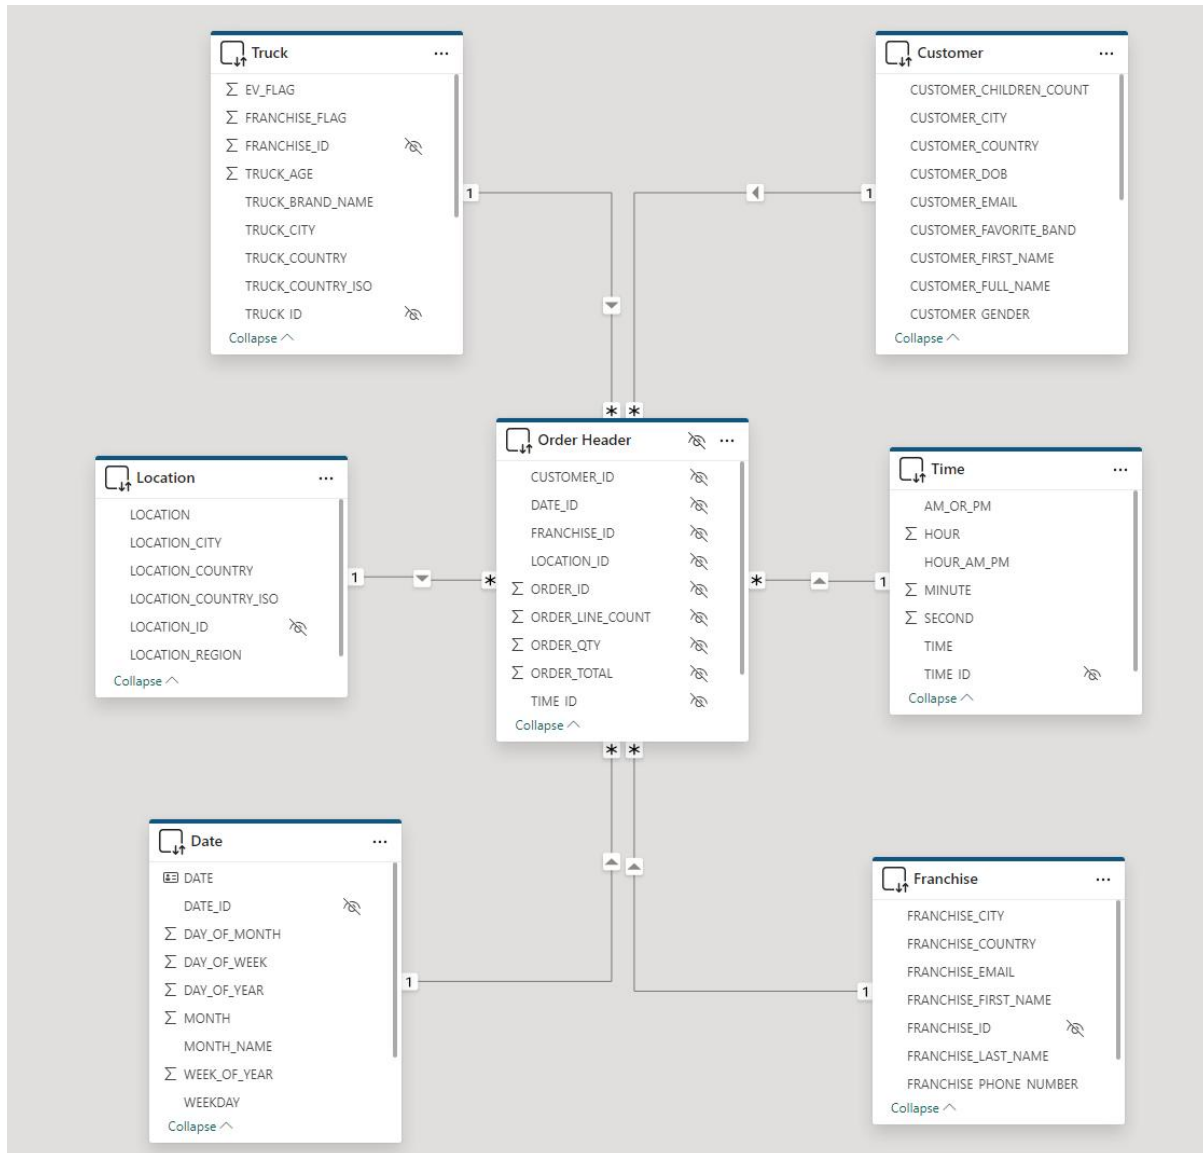


Figure 3: Illustration of a relational schema model, showcasing entity relationships between trucks, customers, locations, orders, time, dates, and franchises for structured data analysis.

While star schemas themselves are platform-agnostic, Snowflake’s approach to building and maintaining them can be greatly simplified using Dynamic Tables. Dynamic Tables are a unique feature in Snowflake that let data teams define SQL transformations in a declarative way. Snowflake then orchestrates the refresh of these tables on a schedule, optionally leveraging incremental updates.

Unlike a standard view, a Dynamic Table persists the results into physical storage, which can lead to significant performance benefits. It does so without requiring manual scheduling of the transformations in external pipelines. Instead, the user simply specifies a **refresh\_mode** (incremental or full), a **target\_lag**, and a **warehouse** to be used for refreshes. Snowflake then executes these transformations according to the given parameters.

In the Tasty Bytes scenario, dimension tables such as **dt\_dim\_truck**, **dt\_dim\_franchise**, **dt\_dim\_menu\_item**, **dt\_dim\_location**, and **dt\_dim\_customer** can each be defined as Dynamic Tables. A fact table, such as **dt\_fact\_order\_detail**, might also be a Dynamic Table. Code blocks for creating these objects appear similar to:

```
create or replace dynamic table dt_dim_truck
  target_lag = 'DOWNSTREAM'
  warehouse = tb_de_wh
  refresh_mode = incremental
  initialize = on_create
as
  select distinct
    t.truck_id,
    t.franchise_id,
    m.truck_brand_name,
    ...
  from tb_101.raw_pos.truck t
  join tb_101.raw_pos.menu m on m.menu_type_id = t.menu_type_id;
```

By referencing the newly created dimension tables inside the definition of a fact table, an entire transformation pipeline can be built with a series of chained Dynamic Tables. In the example code, `dt_fact_order_header` might aggregate the more granular `dt_fact_order_detail` to facilitate high-level reporting. Another Dynamic Table, such as `dt_fact_order_agg`, might further roll up data to reduce cardinalities for certain reporting scenarios.

Once created, these Dynamic Tables are visible in Snowsight, and data engineers can observe their lineage graph in the UI. This lineage graph helps users visualize how each table depends on others, identify bottlenecks, and see how changes propagate. By the end of this transformation phase, the Tasty Bytes data is neatly arranged in a star schema, making it ready for consumption by Power BI.

## VI. PROTECTING SENSITIVE DATA WITH SNOWFLAKE HORIZON

Data governance is a pressing concern for organizations of all sizes. Heightened data privacy regulations such as the GDPR, CCPA, and emerging data protection laws in various countries require robust governance strategies. Snowflake's solution, known as Snowflake Horizon, spans data discovery, governance, and protections. This section highlights how to label sensitive fields, apply masking policies to obfuscate data at query time, and configure row-access policies to ensure analysts only see the data they are entitled to see.

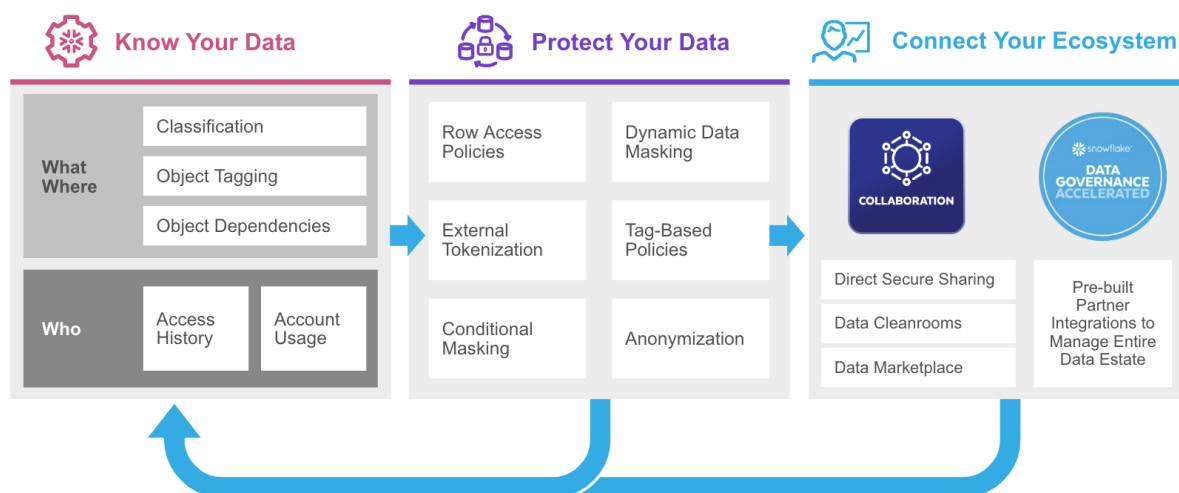




Figure 4: Illustration of Snowflake's data governance framework, emphasizing data classification, protection, and secure ecosystem integration for enhanced compliance and collaboration.

Snowflake allows for both system-defined and user-defined tags. Classification is the system-driven process that scans columns to detect sensitive fields based on patterns (such as phone numbers, email addresses, and personally identifiable information). For instance, a user may run a stored procedure like:

```
call system$classify('tb_101.powerbi.dt_dim_customer', {'auto_tag': true});
```

This procedure instructs Snowflake to analyze the data and automatically apply relevant system tags (for example, `SEMANTIC_CATEGORY.EMAIL` or `SEMANTIC_CATEGORY.PHONE_NUMBER`). A user may further define custom tags for more specific use cases:

```
create or replace tag tb_101.powerbi.pii_name_tag;
...
alter table tb_101.powerbi.dt_dim_customer
  modify column customer_first_name
    set tag tb_101.powerbi.pii_name_tag = 'First Name';
```

Tags help data stewards know which columns contain PII or other sensitive data, thereby guiding subsequent policy application.

Dynamic Data Masking in Snowflake is an essential feature for column-level security. When a masking policy is applied to a specific column, the data might remain in clear-text physically, but queries will retrieve masked values depending on the user's role. For example, if we create a masking policy such as:

```
create or replace masking policy tb_101.powerbi.name_mask AS (val STRING) RETURNS
STRING ->
  case
    when CURRENT_ROLE() IN ('SYSADMIN', 'ACCOUNTADMIN',
'TB_BI_ANALYST_GLOBAL') THEN val
    else '***~MASKED~***'
  end;
```

Then assign it to a tag:

```
alter tag tb_101.powerbi.pii_name_tag
  set masking policy tb_101.powerbi.name_mask;
```

Any column tagged with `tb_101.powerbi.pii_name_tag` becomes automatically masked at query time for roles not explicitly granted viewing privileges. This approach drastically simplifies large-scale governance; the data steward updates one policy and it propagates to all columns carrying the same tag.

In addition to column-level masking, row-level security can be enforced with row-access policies. A typical approach is to create a “mapping table” that defines which roles can see certain subsets of the data. For instance, if Tasty Bytes employs regional analysts, each region might be mapped to location IDs in a table `tb_101.powerbi.row_policy_map`. The row-access policy references this mapping table in its logic, often with an `EXISTS` clause to check if the current role is authorized for that row:

```
create or replace row access policy tb_101.powerbi.rap_dim_location_policy
as (location_id NUMBER) returns boolean ->
CURRENT_ROLE() in
('ACCOUNTADMIN','SYSADMIN', 'TB_BI_ANALYST_GLOBAL')
or exists
(
select rp.role
from tb_101.powerbi.row_policy_map rp
where rp.role = CURRENT_ROLE()
and rp.location_id = location_id
);
```

Once attached to tables containing location IDs (for example, the fact tables and location dimension), analysts using the role `tb_bi_analyst_na` will only see rows pertaining to North American locations. Combined with dynamic masking, this row-level security ensures a robust, multi-layered data governance that is automatically enforced at query time, even when connecting through external BI tools such as Power BI.

## VII. ANALYZING SNOWFLAKE DATA IN POWER BI

With data transformations completed and governance rules in place, the next step is to connect Power BI to the curated star schema in Snowflake. The final result is typically a Power BI dataset (or semantic model) that references each dimension and fact table. This section highlights the steps involved in connecting to Snowflake, designing a semantic model, and verifying that governance rules continue to work seamlessly in Power BI.

In many projects, data teams or solution architects prepare a Power BI Template file (.pbit). A template file stores metadata such as table definitions, relationships, measures, and connections, but does not store the actual data. Upon opening the template in Power BI Desktop, a user can supply parameters for their own Snowflake account, warehouse, database, schema, or role.

In this Tasty Bytes illustration, the user opens a provided .pbit file in Power BI Desktop. A parameters screen might prompt for:

```
Snowflake Account = xy12345.us-east-2.aws.snowflakecomputing.com
Snowflake Warehouse = TB_POWERBI_WH
Snowflake Role = TB_BI_ANALYST_GLOBAL
Snowflake Database = TB_101
Snowflake Schema = POWERBI
```

Once the user supplies these parameters, Power BI prompts for authentication. Organizations that have configured SSO or external OAuth can skip the manual password entry, but for demo or quickstart purposes, a

user password-based authentication might be used. After successfully connecting, the dataset is loaded, with a default page that might include summary metrics like total sales and total orders. Because the dataset is in DirectQuery mode, each visual's data is retrieved by generating SQL queries that Power BI sends to Snowflake. This arrangement allows near real-time data because there is no scheduled refresh required. If the data in Snowflake updates, queries from Power BI will reflect those changes almost immediately.

Once the connection is established, users typically open the model view in Power BI to see the relationships among the dimension and fact tables. Because the star schema in Snowflake has a clear separation of facts and dimensions, the Power BI relationships are usually one-to-many, with the dimension side being “one.” This structure is simpler to maintain than if the data were flattened into a giant denormalized table.

In a star schema scenario, the user might see the following objects:

```
Dim_Date
Dim_Time
dt_dim_customer
dt_dim_franchise
dt_dim_location
dt_dim_truck
dt_dim_menu_item
dt_fact_order_detail
dt_fact_order_header
dt_fact_order_agg
```

Although we are avoiding explicit bullet points in this article, it is useful to conceptually group them as dimensions or facts. Each dimension has a key column that ties to the foreign key in the fact tables, forming the star.

In Power BI, many designers mark dimension tables as “Dimension” or “Lookup” roles and fact tables as “Fact.” Sometimes, measures such as total sales, total orders, or average unit price are defined in a measure table or directly in the fact tables. Because this scenario includes an aggregated fact table (`dt_fact_order_agg`), the data model in Power BI can incorporate user-defined aggregations, directing queries to the smaller `dt_fact_order_agg` whenever possible and falling back to more granular tables if needed. This approach avoids scanning hundreds of millions of rows each time a user only needs aggregated information.

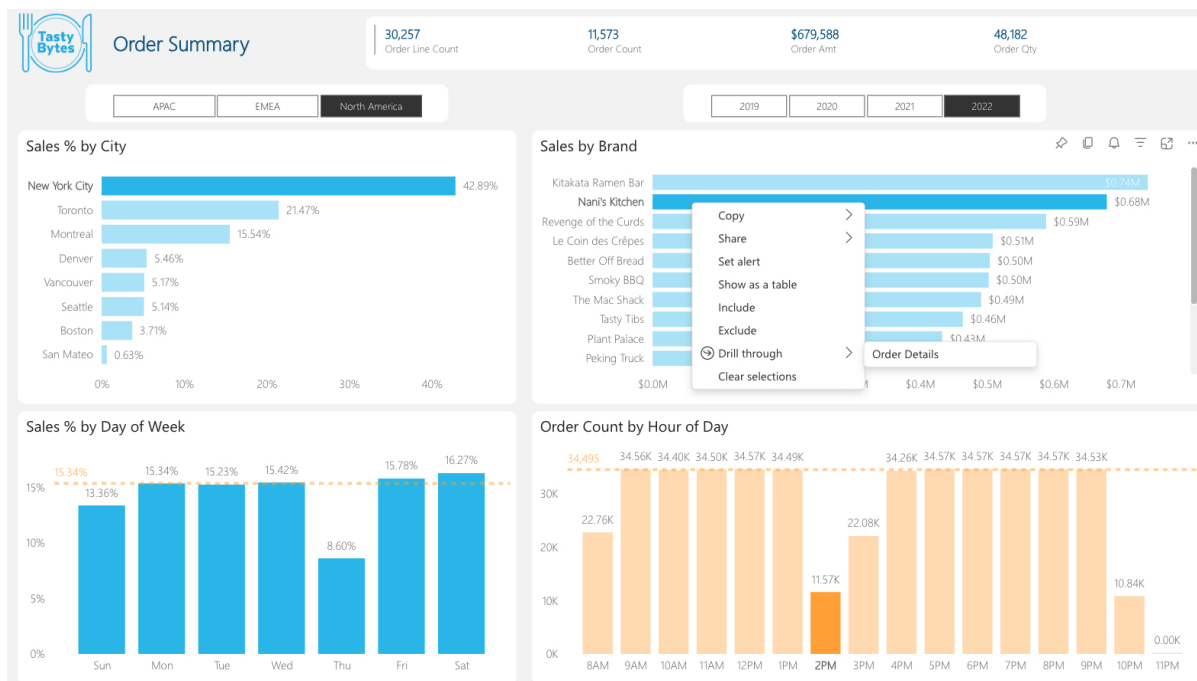


Figure 5: Illustration of exploring Tasty Bytes data via Power BI

Once the dataset is validated, it can be published to the Power BI Service, typically to a workspace designated for data analytics. Publishing creates two artifacts: a semantic model (the dataset) and a report. In the workspace settings, the user can modify connection parameters, such as changing the Snowflake role from **TB\_BI\_ANALYST\_GLOBAL** to another role. Because the model is in DirectQuery mode, there is no need to schedule a refresh. The queries will be run on-demand whenever a user interacts with the report.

A major reason to keep transformations within Snowflake is that the governance and security measures configured in Snowflake remain in force. When a user selects a different Snowflake role in the Power BI dataset parameters, the queries from Power BI run as that role, triggering the correct masking or row-access policies. For instance, if the role is changed to **TB\_BI\_ANALYST\_NA**, the user might see masked customer names (depending on the policy) and only rows that correspond to North America. If the user changes the parameter to **TB\_BI\_ANALYST\_EMEA**, then the same visuals or tables in the report show only European data. This dynamic security reduces the complexity of implementing equivalent row-level security in Power BI and ensures consistent data governance across all entry points to Snowflake.

In a typical demonstration, the user might navigate to a page with a table or matrix visual that includes customer email addresses or phone numbers. If the user is connected as **TB\_BI\_ANALYST\_GLOBAL**, the table might show real values. If the connection parameter changes to **TB\_BI\_ANALYST\_NA** or **TB\_BI\_ANALYST\_APAC**, the same table might display partially masked phone numbers or email addresses replaced with **\*\*\*\*\*@example.com**. Such a scenario proves that the policies configured earlier in Snowflake function as intended.

## VIII. CONCLUSION AND FUTURE DIRECTIONS

The combination of Snowflake and Power BI has become a de facto standard for numerous organizations seeking cloud-scale analytics. This article illustrated an end-to-end analytics workflow that begins with raw data profiling in Snowflake, continues with the creation of roles and warehouses for dedicated BI usage, includes the

enrichment of data with third-party information from the Snowflake Marketplace, transforms the data using Dynamic Tables organized into a star schema, applies robust data masking and row-level security policies with Snowflake Horizon, and culminates with a near real-time analytics experience in Power BI Desktop or the Power BI Service.

Several foundational concepts, including dimensional modeling, role-based access control, star schemas, and dynamic data masking, remain at the heart of modern data warehousing. However, new Snowflake features such as Dynamic Tables for declarative pipelines and the Snowflake Marketplace for frictionless third-party data enrichments dramatically reduce the time to value. On the Power BI side, the improvements in DirectQuery performance, aggregated table functionality, and advanced modeling features allow large data volumes in Snowflake to be explored in near real-time, avoiding the constraints imposed by in-memory data caches.

In a research context, the success of this integration can be measured in multiple ways. Performance metrics can be collected to see how quickly visuals render under increasing user concurrency. Security tests can confirm that row-access policies are enforced when roles change. User adoption metrics can indicate how effectively data consumers across different lines of business can self-serve analytics without repeatedly relying on data engineers to modify transformations. As data volumes increase and analytics needs become more varied, both Snowflake and Power BI will continue to enhance features that expand flexibility, performance, and security.

Future directions for this integrated approach may include deeper machine learning and AI workflows on the Snowflake side, with the results exposed in Power BI or integrated into advanced analytics tools. There is also ongoing development in the area of external OAuth and SSO-based authentication, allowing streamlined identity federation so that Power BI's user context is passed all the way into Snowflake. Additionally, organizations may explore more sophisticated real-time streaming scenarios, hooking up event-driven architectures that feed into Snowflake and appear almost immediately in Power BI dashboards.

Regardless of the future's technical developments, the fundamental best practices remain. Teams must establish proper data governance, ensure that the data transformations happen upstream rather than in scattered logic throughout the BI tool, and employ an architecture that scales elastically to handle both sporadic bursts and steady production loads. With correct planning, Snowflake's cloud data platform and Power BI's intuitive analytics layer can form a powerful, future-proof analytics solution capable of accommodating the most demanding business scenarios.

## IX. REFERENCES

- [1] N. Kumar, "Big Data Analytics in Cloud – Comparative Study," *Ijcttjournal.org*, 2023.
- [2] Bhavik Merchant; Christopher Webb, "Microsoft Power BI Performance Best Practices: A comprehensive guide to building consistently fast Power BI solutions," Packt Publishing, 2022
- [3] M. Jaipurkar, N. Ragit, C. Tambuskar, and P. D. Saraf, "IPL Data Analysis and Visualization Using Microsoft Power BI Tool," 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT), pp. 1–6, Jul. 2023.
- [4] Rinkesh Gajera, "Integrating Power Bi with Project Control Systems: Enhancing Real-Time Cost Tracking and Visualization in Construction," *International Journal of Scientific Research in Civil Engineering*, vol. 7, no. 5, pp. 154–160, Oct. 2023.

- [5] K. Kishor, N. A. Joshi, P. Singh, None Akshun Chhapola, N. S. Jain, and A. Gupta, “Leveraging Power BI for Enhanced Data Visualization and Business Intelligence,” Universal Research Reports, vol. 10, no. 2, pp. 676–711, 2023.
- [6] “End-to-End Analytics with Snowflake and Power BI,” Snowflake.com, 2018.
- [7] Becker, L. T., & Gould, E. M. (2019). Microsoft Power BI: Extending excel to manipulate, analyze, and visualize diverse data. Serials Review, 45(3), 184–188, 2019.
- [8] Serge Gershkovich, Kent Graziano, “Data Modeling with Snowflake: A practical guide to accelerating Snowflake development using universal data modeling techniques,” Packt Publishing, 2023.
- [9] F. Bell, Raj Chirumamilla, B. B. Joshi, B. Lindstrom, R. Soni, and Sameer Videkar, “Data Sharing, Data Exchanges, and the Snowflake Data Marketplace,” Apress eBooks, pp. 299–328, Dec. 2021.
- [10] Dr. Sunil. Khilari, C. Singh, and Mr. Bharat. Mane, “Business Intelligence Tool-Power BI for Performance Management,” SSRN Electronic Journal, 2022.