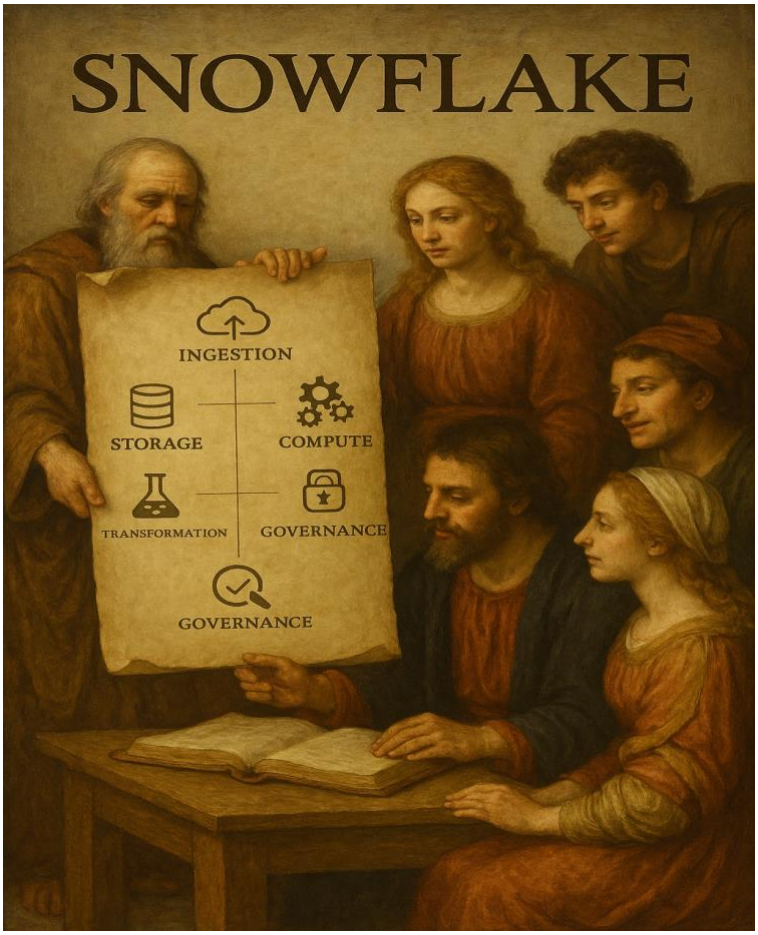

SNOWFLAKE

A brief scroll — Think less “treatise,” more “continental cafe chat with a diagram.”

ABSTRACT

Outlines the Snowflake architecture in practical terms—how it ingests, stores, transforms, secures, and serves data. Whether you're feeding dashboards, automating reports, or publishing data products, Snowflake sits at the centre, blending performance with governance. The following sections map out the flow and explain each building block using a shorthand delivery-focused narrative.

.

Iain Hamilton Toolin

1.2 Architecture of Pipeline (Snowflake Blackbox)

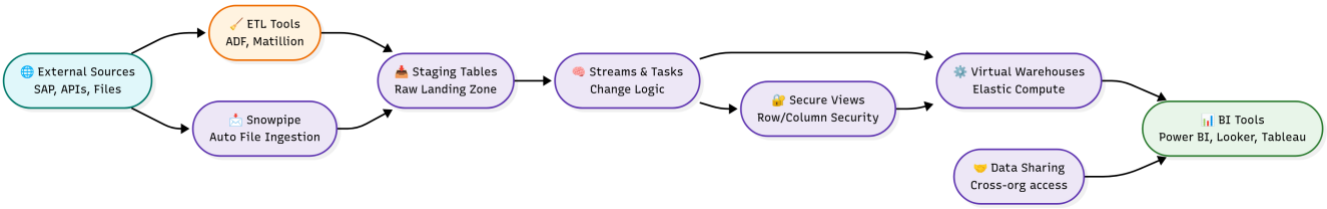

Figure 1: Snowflake pipeline architecture

1.3 Each step in the pipeline is described using plain English and contextual usage.

Component	What It Is	Narrative
External Sources	SAP, Salesforce, flat files, APIs	This is your raw feed—fragmented and full of potential.
ETL Tools (ADF, Matillion)	Transformation before or within Snowflake	The scrubbing squad—pulling in the right data at the right time.
Snowpipe	Automated file ingestion	Drop a file, and Snowpipe does the rest—no manual triggers needed.
Staging Tables	Landing zone for raw data	Like an inbox—temporary, untidy, but essential.
Streams and Tasks	Change tracking and scheduled logic	How data gets refined—tracked, versioned, and made useful.
Secure Views	Governed access layer	Sliced and filtered views of data, with row-level security and masking.
Data Sharing	Direct share with partners or consumers	No duplication—just permissioned access across clouds or orgs.

Virtual Warehouses	Elastic compute for querying	Your on-demand muscle—spin up, crunch, shut down.
BI / Analytics Tools	Power BI, Tableau, Looker	This is the visible layer—where all the behind-the-scenes effort pays off.

Table 1: Narrative on services and roles.

1.4 Snowflake: Data Journey Pipeline (white box)

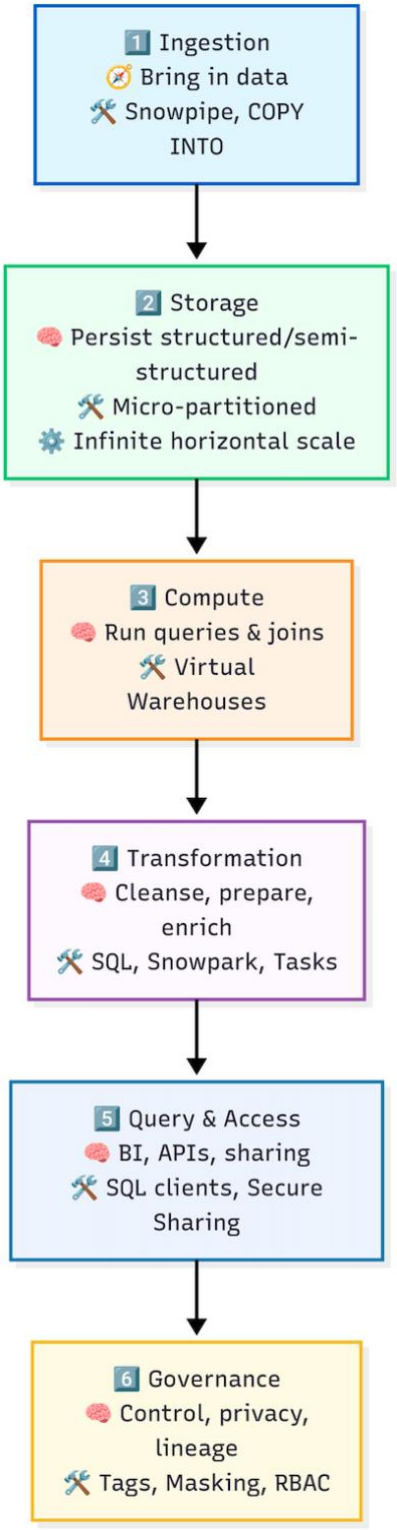

Figure 2: Snowflake pipeline architecture

*The intellectual property for the Snowflake platform and its associated architecture belongs to **Snowflake Inc.**, the company that owns and develops the technology. The accompanying narrative and visual interpretation presented here are based on independent research and contextualisation carried out by **Iain Toolin**, and are intended to aid understanding through a pedagogical and stylistic lens. This work is not affiliated with or endorsed by Snowflake Inc*

1.5 Snowflake: Data Journey Table

Step	Stage	Purpose	Tools / Features	Scalability Mechanism
	Ingestion	Bring in data from files, events, APIs	Snowpipe, COPY INTO, Kafka Connect, External Tables	External to Snowflake, but decoupled and parallel-capable
	Storage	Persist structured & semi-structured data	Micro-partitioned columnar storage on S3/Blob/GCS	Unlimited horizontal scale — managed blob storage
	Compute	Execute queries, transforms, joins	Virtual Warehouses (size & multi-cluster)	✅ **Vertical:** resize warehouse ✅ **Horizontal:** scale-out multi-cluster
	Transformation	Data prep, cleansing, ML features	SQL, Snowpark, Streams, Tasks, UDFs, External Functions	Compute-driven: parallelism and independent scheduling
	Query & Access	Data consumption by users, apps, BI	SQL clients, BI tools, APIs, Secure Data Sharing	Independent compute clusters + result caching
	Governance	Metadata, lineage, access control, privacy	Tags, masking policies, RBAC, Access History	Centralised control – no scaling bottlenecks

Table 2: Narrative on Snowflake services.