

Data Warehousing

Reference: Data Mining: Concepts and Techniques by Jiawei Han and Micheline Kamber

Data Warehousing: Purpose and Strategic Value

Modern organizations operate in environments where competition is intense, customer expectations change rapidly, and decisions must be made based on evidence rather than intuition. Data warehousing emerged to address this reality by providing a systematic way to organize, integrate, and analyze large volumes of data accumulated across an enterprise.

At its core, a data warehouse is not merely a database. It is a decision-support infrastructure that consolidates data from multiple operational systems and external sources, restructures it for analysis, and preserves historical context. This enables managers, analysts, and executives to identify trends, evaluate performance, and make informed strategic decisions. Over time, many organizations have recognized that understanding customer behavior, operational efficiency, and market dynamics through integrated data is a competitive necessity rather than a luxury.

What a Data Warehouse Is—and What It Is Not

A data warehouse is maintained separately from operational databases. Operational systems are designed to record day-to-day transactions—sales, payments, registrations, updates—efficiently and accurately. A data warehouse, by contrast, is designed for analysis rather than transaction processing.

A widely accepted definition describes a data warehouse as a collection of data that is:

- **Subject-oriented:** Data is organized around key business subjects such as customers, products, sales, or suppliers. This orientation simplifies analysis by focusing on what decision makers care about, rather than on operational processes.
- **Integrated:** Data from heterogeneous sources is brought together into a consistent format. Differences in naming conventions, data types, units, and codes are resolved through data cleaning and integration.
- **Time-variant:** The warehouse stores historical data, often spanning several years. Each record is associated with a time dimension, enabling trend analysis and longitudinal studies.
- **Nonvolatile:** Once data is loaded, it is not updated or deleted in the same way as operational data. The warehouse is stable, supporting read-intensive analytical queries rather than frequent insert, update, or delete operations.

Together, these characteristics distinguish a data warehouse from traditional databases and make it suitable for decision support.

How Organizations Use Data Warehouses

Organizations rely on data warehouses to support a wide range of analytical and strategic activities, including:

- Understanding customer behavior by analyzing buying patterns, preferences, timing, and spending habits.
- Evaluating product performance across time periods and geographic regions to refine pricing, production, and distribution strategies.
- Identifying inefficiencies or profit opportunities in operations.
- Supporting customer relationship management, asset utilization, cost control, and environmental or regulatory analysis.

An additional and often underappreciated benefit is the role of data warehousing in integrating heterogeneous databases. Instead of querying multiple autonomous systems on demand—a process that is complex, slow, and resource-intensive—data from these systems is integrated in advance. This results in faster queries, reduced interference with operational workloads, and a unified semantic view of enterprise data.

Operational Systems Versus Analytical Systems

To fully appreciate the role of a data warehouse, it is useful to contrast operational database systems with analytical systems.


Operational systems are designed for transaction processing. They support large numbers of concurrent users, ensure data consistency through locking and logging, and focus on current, highly detailed data. Their performance is measured in terms of transaction throughput and availability.


Analytical systems, often referred to as online analytical processing systems, are designed for exploration and analysis. They manage large volumes of historical data, support aggregation and summarization, and allow users to view data at multiple levels of detail. Queries are typically complex and read-intensive, and performance is measured by query response time and flexibility.

Running analytical queries directly on operational systems would severely degrade transactional performance and complicate concurrency control. For this reason, maintaining a separate data warehouse remains the dominant architectural approach, even though modern database systems increasingly blur the boundary between the two.

Multitier Architecture of a Data Warehouse

Data warehouses are commonly implemented using a multitier architecture that separates data storage, analytical processing, and user interaction.





At the foundation lies the warehouse database server, typically implemented using relational database technology. Data is extracted from operational systems and external sources, cleaned to remove errors and inconsistencies, transformed into a unified format, and loaded into the warehouse. This layer also maintains metadata describing the structure, content, and lineage of the data.

Above this sits the analytical server layer. This layer supports multidimensional analysis and may be implemented using relational extensions or specialized multidimensional storage. It translates analytical requests into efficient data access paths.

At the top are client-facing tools used by analysts and decision makers. These include query and reporting tools, dashboards, visualization systems, and data mining or predictive analytics tools.

Enterprise Warehouses, Data Marts, and Virtual Warehouses

From an architectural perspective, data warehouses can be organized in several ways.

An **enterprise warehouse** integrates data across the entire organization. It provides a comprehensive, consistent view of enterprise data and typically contains both detailed and summarized information. While powerful, it is expensive and time-consuming to build.

A **data mart** is a focused subset of data designed for a specific department or business function, such as marketing or finance. Data marts are faster to implement and less costly, but if designed in isolation they can lead to integration challenges later.

A **virtual warehouse** consists of views defined over operational databases. It is easy to create but places additional load on operational systems and offers limited performance for complex queries.

In practice, many organizations adopt an incremental approach: defining a high-level corporate data model early, building data marts in parallel, and gradually integrating them into a cohesive enterprise warehouse.

Extraction, Transformation, and Loading

Populating a data warehouse requires a disciplined backend process commonly referred to as extraction, transformation, and loading.

- **Extraction** collects data from diverse internal and external sources.
- **Cleaning** identifies and corrects errors, inconsistencies, and missing values.
- **Transformation** converts data into formats and structures suitable for analytical use.
- **Loading** inserts data into the warehouse, often after sorting, summarizing, and indexing.
- **Refresh** propagates updates from source systems into the warehouse at scheduled intervals.

The quality of these processes directly affects the reliability of analytical results.

The Role of Metadata

Metadata is often described as “data about data,” but in a data warehouse it plays a central operational and analytical role. Metadata documents the structure of the warehouse, the meaning of data elements, data lineage, transformation rules, aggregation logic, refresh schedules, and access controls.

For analysts, metadata serves as a directory that explains what data exists and how it can be used. For system designers, it provides the blueprint that connects operational data to analytical representations. Because of its importance, metadata is stored persistently and managed with the same care as the warehouse data itself.

Multidimensional Modeling and the Data Cube

Analytical processing in a data warehouse is based on a multidimensional view of data, commonly conceptualized as a data cube.

2-D View of Sales Data for *AllElectronics* According to *time* and *item*

		<i>location</i> = "Vancouver"				
		<i>item</i> (type)				
<i>time</i> (quarter)		<i>home</i>				
		<i>entertainment</i>	<i>computer</i>	<i>phone</i>	<i>security</i>	
Q1		605	825	14	400	
Q2		680	952	31	512	
Q3		812	1023	30	501	
Q4		927	1038	38	580	

Note: The sales are from branches located in the city of Vancouver. The measure displayed is *dollars_sold* (in thousands).


3-D View of Sales Data for *AllElectronics* According to *time*, *item*, and *location*

<i>location</i> = "Chicago"					<i>location</i> = "New York"					<i>location</i> = "Toronto"					<i>location</i> = "Vancouver"				
<i>Item</i>					<i>Item</i>					<i>Item</i>					<i>Item</i>				
<i>home</i>					<i>home</i>					<i>home</i>					<i>home</i>				
<i>time</i>	<i>ent.</i>	<i>comp.</i>	<i>phone</i>	<i>sec.</i>	<i>ent.</i>	<i>comp.</i>	<i>phone</i>	<i>sec.</i>		<i>ent.</i>	<i>comp.</i>	<i>phone</i>	<i>sec.</i>		<i>ent.</i>	<i>comp.</i>	<i>phone</i>	<i>sec.</i>	
Q1	854	882	89	623	1087	968	38	872	818	746	43	591	605	825	14	400			
Q2	943	890	64	698	1130	1024	41	925	894	769	52	682	680	952	31	512			
Q3	1032	924	59	789	1034	1048	45	1002	940	795	58	728	812	1023	30	501			
Q4	1129	992	63	870	1142	1091	54	984	978	864	59	784	927	1038	38	580			

Note: The measure displayed is *dollars_sold* (in thousands).

location (cities)		Chicago				New York				Toronto				Vancouver			
		854	882	89	623	1087	968	38	872	818	746	43	591	698	682	925	789
time (quarters)	Q1	605	825	14	400												
	Q2	680	952	31	512										728	1002	870
	Q3	812	1023	30	501										784	984	
	Q4	927	1038	38	580												
		computer	security														
		home	phone														
		entertainment															
		item (types)															

A 3-D data cube representation of the data in Table 4.3, according to *time*, *item*, and *location*. The measure displayed is *dollars_sold* (in thousands).




A 4-D data cube representation of sales data, according to *time*, *item*, *location*, and *supplier*. The measure displayed is *dollars_sold* (in thousands). For improved readability, only some of the cube values are shown.

In this model, **dimensions** represent perspectives such as time, location, product, or customer. **Facts** are numeric measures, such as sales amount or units sold, that quantify business activity. A fact table stores these measures along with keys referencing the associated dimension tables.

Although the term “cube” suggests three dimensions, real-world data cubes are n-dimensional. They support analysis at multiple levels of aggregation. The most detailed representation is known as the base cuboid, while the most summarized representation, aggregating across all dimensions, is called the apex cuboid. Together, all possible aggregations form a lattice of cuboids that constitutes the full data cube.


Schema Designs for Multidimensional Data

To implement multidimensional models in practice, several schema designs are commonly used.




Star schema of *sales* data warehouse.

A **star schema** consists of a central fact table connected to denormalized dimension tables. It is simple, intuitive, and efficient for querying, making it the most widely used design. Star schema: The most common modeling paradigm is the star schema, in which the data warehouse contains (1) a large central table (fact table) containing the bulk of the data, with no redundancy, and (2) a set of smaller attendant tables (dimension tables), one for each dimension. The schema graph resembles a starburst, with the dimension tables displayed in a radial pattern around the central fact table.



Snowflake schema of *sales* data warehouse.

A **snowflake schema** normalizes some dimension tables into multiple related tables. This reduces redundancy but increases query complexity due to additional joins.




Fact constellation schema of a sales and shipping data warehouse.

A **fact constellation** schema supports multiple fact tables that share common dimensions. This design is typical in enterprise warehouses that model multiple, related business processes.

What is OLAP?

- **OLAP (Online Analytical Processing)** is a technology for **fast, interactive analysis** of large volumes of **multidimensional data**.
- It supports analysis across dimensions such as **time, location, product, customer, etc.**
- OLAP is designed for **decision support**, not routine transactions (unlike OLTP).
- It enables **complex queries, aggregations, trend analysis, and what-if analysis**.
- Core OLAP operations include **roll-up, drill-down, slice, dice, and pivot**.



Examples of typical OLAP operations on multidimensional data.

Concept Hierarchies in OLAP

- In a **multidimensional model**, each **dimension** has multiple **levels of abstraction**.
- These levels are organized as **concept hierarchies**.
- Concept hierarchies allow users to:
 - View data at **different granularities**
 - **Summarize** or **detail** data dynamically
- Example hierarchies:
 - **Time:** day → month → quarter → year
 - **Location:** street → city → state/province → country
 - **Item:** item → item type → category

Key Benefit: Flexible, user-friendly **interactive data analysis** from multiple perspectives.

Example OLAP Data Cube

- **Dimensions:**
 - **Location** (city level)
 - **Time** (quarter level)
 - **Item** (item type level)
- **Measure:**

- Sales amount (dollars sold, in thousands)
- Sample cities analyzed:
 - Chicago, New York, Toronto, Vancouver
- This cube is referred to as the **central cube**.

Typical OLAP Operations

Roll-up (Drill-up)

Purpose: Aggregate data to a **higher level**.

- Performed by:
 - Climbing up a concept hierarchy, or
 - Reducing dimensions
- Example (Hierarchy-based):
 - Location: **city** → **country**
 - Result: Sales grouped by **country**, not city
- Example (Dimension reduction):
 - Removing **time** dimension
 - Result: Total sales by **location only**

Drill-down

Purpose: Move from **summary data to detailed data**.

- Reverse of roll-up
- Performed by:
 - Descending a concept hierarchy, or
 - Adding new dimensions
- Example (Hierarchy-based):
 - Time: **quarter** → **month**
 - Result: Monthly sales instead of quarterly
- Example (Dimension addition):
 - Adding **customer group** dimension for finer analysis

Slice

Purpose: Select data on **one dimension**.

- Fixes a single dimension value
- Result: A **subcube**
- Example:
 - Time = **Q1**

- Result: Sales data only for Quarter 1

Dice

Purpose: Select data on **multiple dimensions**.

- Applies multiple filtering conditions
- Result: A **smaller subcube**
- Example:
 - Location = Toronto or Vancouver
 - Time = Q1 or Q2
 - Item = Home entertainment or Computer

Pivot (Rotate)

Purpose: Change **data orientation** for better visualization.

- Rotates axes of the cube
- Does **not change data**, only presentation
- Examples:
 - Swap **item** and **location** axes
 - Convert a 3D cube into multiple 2D views

Other OLAP Operations

- **Drill-across:**
 - Queries across **multiple fact tables**
- **Drill-through:**
 - Accesses **underlying relational tables** using SQL
- **Advanced analytics:**
 - Top-N / Bottom-N ranking
 - Moving averages
 - Growth rates
 - Currency conversion
 - Depreciation, interest, IRR
 - Statistical functions

Analytical Power of OLAP

- Built-in **calculation engine** for:
 - Ratios, variance, derived measures

- Supports:
 - Aggregations at **all granularities**
 - Analysis across **all dimension intersections**
- Enables:
 - **Forecasting**
 - **Trend analysis**
 - **Statistical modeling**

Conclusion: OLAP engines are powerful tools for **business intelligence and strategic analysis**.

OLAP Systems vs Statistical Databases (SDBs)

Similarities:

- Multidimensional data models
- Concept hierarchies
- Measures linked with dimensions
- Roll-up and drill-down operations

Difference:

- OLAP evolved for **business decision support**
- SDBs focus on **statistical applications**
- Differences largely arise from **terminology and application domains**

Starnet Query Model (Multidimensional Querying)

What is a Starnet Model?

- A **visual query model** for OLAP
- Consists of:
 - A **central point** (fact table)
 - **Radial lines** (dimensions)
- Each radial line represents a **concept hierarchy**
- Each level in a hierarchy is called a **footprint**

Example: AllElectronics Starnet

- **Dimensions (Radial Lines):**
 - Location
 - Customer
 - Item

- Time
- **Time footprints:**
 - Day → Month → Quarter → Year
- **Location hierarchy:**
 - Street → City → State/Province → Country

Role of Concept Hierarchies in Starnet

- Enable:
 - **Generalization** (day → year)
 - **Specialization** (country → city)
- Support OLAP operations:
 - **Roll-up**
 - **Drill-down**
- Allow users to analyze data at **any desired granularity**

OLAP combines multidimensional modeling, concept hierarchies, and powerful operations to enable fast, flexible, and interactive decision-oriented data analysis.

Part-wise Explanation of the above figure(Fig-OLAP OPERATIONS)

Central Cube (Middle of the Figure)

- **Referenced part:** The large cube located at the center
- **Meaning:** This is the **central OLAP data cube** for *AllElectronics sales*.
- **Dimensions shown:**
 - **Location** □ Chicago, New York, Toronto, Vancouver
 - **Time** □ Q1, Q2, Q3, Q4
 - **Item** □ computer, phone, home entertainment, security
- **Measure:** Sales amount (dollars sold).

Roll-up Operation (Upper-Right Cube)

- **Referenced part:** Cube labeled □roll-up (from cities to country)□
- **Meaning:** Sales data are **aggregated from city level to country level (USA)**.
- **Concept hierarchy used:** city □ country
- **Effect:** Fewer cells, higher-level summarized view.

Drill-down Operation (Right-Side Tall Cube)

- **Referenced part:** Cube labeled □drill-down (from quarters to months)□
- **Meaning:** Time dimension is expanded from **quarters to individual months**.
- **Concept hierarchy used:** quarter □ month
- **Effect:** More detailed, fine-grained sales information.

Slice and Dice Operations (Top-Left and Middle-Left Cubes)

- **Slice**
 - **Referenced part:** *Top-left cube labeled \square slice (time = Q1)* \square
 - **Meaning:** Only **Q1 data** is selected, reducing the cube to a subcube.
- **Dice**
 - **Referenced part:** *Middle-left cube with multiple selection conditions*
 - **Meaning:** Filters applied on **location, time, and item** simultaneously to create a smaller subcube.

Pivot (Bottom-Left Table)

- **Referenced part:** *Bottom-left 2-D table labeled \square pivot* \square
- **Meaning:** Axes are **rotated** (item vs location) to change the **visual orientation**.
- **Effect:** Data values remain unchanged, but the **viewpoint is altered** for better analysis.

>>Each region of above figure demonstrates a **specific OLAP operation**--central cube (base data), upper-right (roll-up), right-side (drill-down), left-side (slice/dice), and bottom-left (pivot)--showing how multidimensional data can be interactively analyzed.

Querying Multidimensional Databases with the Starnet Model

Querying multidimensional data is most effective when users can naturally navigate between different levels of abstraction. The **starnet query model** provides an intuitive framework for this navigation.

In the starnet model, queries are visualized as **radial lines extending from a central point**, where each radial line represents a **dimension** of analysis such as time, location, item, or customer. Along each line are **footprints**, which correspond to abstraction levels within a concept hierarchy. These footprints define the granularities available for analysis.




Fig- A Starnet model of the business queries.

For example, the time dimension may include footprints such as day, month, quarter, and year. The location dimension may progress from street to city, province or state, and country. By moving upward or downward along these radial lines, users perform **roll-up** and **drill-down** operations. Rolling up replaces detailed data with more generalized summaries, while drilling down reveals finer detail.

The starnet model is especially useful because it mirrors how analysts think: they begin with a high-level overview and then selectively refine their focus, or they start from detail and gradually abstract patterns.

Why a Business-Driven Design Perspective Matters

A data warehouse is valuable only if it aligns with business objectives. Designing one therefore requires a **business analysis framework** that balances strategic vision with technical feasibility.

Organizations benefit from data warehouses in several fundamental ways. They gain competitive advantage by measuring performance accurately, improve productivity by reducing the time required to gather and reconcile information, strengthen customer relationship management through consistent enterprise-wide views, and reduce costs by identifying long-term trends, anomalies, and inefficiencies.

Designing such a system requires viewing it from multiple complementary perspectives:

- **A business perspective**, which identifies what information is needed now and in the future.
- **A data source perspective**, which reveals what data is currently captured by operational systems and in what form.

- A **warehouse perspective**, which defines how facts, dimensions, and historical context are stored.
- A **business query perspective**, which reflects how end users expect to explore and analyze the data.

These perspectives together ensure that the warehouse serves real decision-making needs rather than becoming a purely technical artifact.

Skills Required to Build and Use a Data Warehouse

Data warehousing is inherently interdisciplinary. Successful implementation requires:

- **Business understanding**, to interpret data meaningfully and translate requirements into analytical questions.
- **Technical expertise**, to design schemas, build extraction and refresh pipelines, and ensure scalability.
- **Analytical skills**, to detect patterns, trends, anomalies, and shifts in behavior.
- **Program management capability**, to coordinate tools, vendors, timelines, and stakeholders.

Without this combination, even technically sound warehouses may fail to deliver value.

Approaches to Data Warehouse Design

Data warehouses can be built using different development strategies.

A **top-down approach** begins with enterprise-wide planning and modeling. It is well suited when business processes are stable and well understood, but it is expensive and slow.

A **bottom-up approach** starts with smaller prototypes or data marts. It delivers faster results and lower initial cost but can create integration challenges if pursued in isolation.

A **combined approach** leverages both: a high-level corporate data model provides consistency, while incremental data marts deliver early benefits and adaptability.

From a development methodology standpoint, iterative and evolutionary methods are often preferred. Rapid cycles allow early validation, continuous refinement, and timely adaptation to changing requirements.

Core Steps in Data Warehouse Design

At a practical level, warehouse design typically follows a sequence of decisions:

1. **Select the business process** to be modeled, such as sales, inventory, or shipments.
2. **Define the grain**, which determines the level of detail stored in the fact table.
3. **Identify dimensions**, such as time, product, customer, or location.
4. **Choose measures**, which are the numeric quantities to be analyzed, such as revenue or units sold.

Clear scoping is critical. Initial implementations should have specific, measurable goals and well-defined constraints on time, budget, and organizational coverage.

How Data Warehouses Are Used Over Time

The use of a data warehouse typically evolves through stages.

Initially, it supports **reporting and predefined queries**. Over time, users begin to explore summarized and detailed data interactively, using charts and dashboards. Later, the warehouse becomes a strategic tool for **multidimensional analysis**, enabling slice-and-dice, pivoting, and trend analysis. Ultimately, it serves as a foundation for **knowledge discovery**, where data mining techniques are applied to uncover deeper insights.

From a functional standpoint, warehouse usage falls into three categories:

- **Information processing**, focused on querying, reporting, and basic statistics.
- **Analytical processing**, focused on OLAP operations across multiple dimensions.
- **Data mining**, focused on discovering hidden patterns, relationships, and predictive models.

OLAP and Data Mining: Complementary Roles

OLAP and data mining are closely related but fundamentally different in purpose.

OLAP emphasizes **user-directed exploration**, summarization, and comparison. It helps analysts understand what has happened and how measures vary across dimensions.

Data mining emphasizes **automated discovery**. It identifies associations, classifications, clusters, and predictive relationships that are not obvious from simple aggregation.

While OLAP simplifies analysis and prepares data at multiple abstraction levels, data mining goes further by revealing implicit knowledge. When combined, they form a powerful analytical environment in which users can interactively guide deeper discovery.

Multidimensional Data Mining

Integrating OLAP with data mining leads to **multidimensional data mining**, sometimes called exploratory or online analytical mining. This integration is particularly powerful because data warehouses already contain clean, integrated, and historical data.

Users can navigate data cubes, select relevant subsets, adjust abstraction levels, and dynamically invoke different mining functions. Visualization tools further enhance understanding by presenting results in intuitive forms.

This integration supports a human-centered approach, where analysts interact with the system rather than relying solely on fully automated discovery.


Efficient Implementation of Data Warehouses


Because data warehouses store massive volumes of data, performance is a central concern. Analytical queries are expected to return results in seconds, even when operating over years of historical data.

Achieving this requires efficient **data cube computation**, indexing, and query processing strategies.

Data Cubes, Group-By Operations, and Dimensional Explosion

A data cube can be viewed as a **lattice of cuboids**, where each cuboid represents a group-by over a specific subset of dimensions.





For a cube with n dimensions, there are potentially 2^n cuboids. When dimensions have multiple hierarchy levels, the number of possible cuboids grows even faster. This rapid growth is known as the **curse of dimensionality**.

Precomputing all cuboids is usually impractical due to storage and maintenance costs.

Materialization Strategies for Data Cubes

There are three basic strategies for cube materialization:

- **No materialization**, where all aggregates are computed on demand.
- **Full materialization**, where all cuboids are precomputed.
- **Partial materialization**, where only selected cuboids or subsets of cells are stored.

Partial materialization offers a balance between storage cost and query performance. Common techniques include storing frequently accessed cuboids, iceberg cubes that retain only significant aggregates, and shell cubes that limit materialization to a subset of dimensions.

Indexing OLAP Data for Performance

Indexing is critical for fast query execution in data warehouses.

Bitmap Indexing

Bitmap indexing represents attribute values using bit vectors, making it highly efficient for low-cardinality dimensions.

Bitmap indexing. In the *AllElectronics* data warehouse, suppose the dimension *item* at the top level has four values (representing item types): “*home entertainment*,” “*computer*,” “*phone*,” and “*security*.” Each value (e.g., “*computer*”) is represented by a bit vector in the *item* bitmap index table. Suppose that the cube is stored as a relation table with 100,000 rows. Because the domain of *item* consists of four values, the bitmap index table requires four bit vectors (or lists), each with 100,000 bits. Figure 4.15 shows a base (data) table containing the dimensions *item* and *city*, and its mapping to bitmap index tables for each of the dimensions. ■

Base table			item bitmap index table					city bitmap index table		
<i>RID</i>	<i>item</i>	<i>city</i>	<i>RID</i>	H	C	P	S	<i>RID</i>	V	T
R1	H	V	R1	1	0	0	0	R1	1	0
R2	C	V	R2	0	1	0	0	R2	1	0
R3	P	V	R3	0	0	1	0	R3	1	0
R4	S	V	R4	0	0	0	1	R4	1	0
R5	H	T	R5	1	0	0	0	R5	0	1
R6	C	T	R6	0	1	0	0	R6	0	1
R7	P	T	R7	0	0	1	0	R7	0	1
R8	S	T	R8	0	0	0	1	R8	0	1


Note: H for “*home entertainment*,” C for “*computer*,” P for “*phone*,” S for “*security*,” V for “*Vancouver*,” T for “*Toronto*.”

Fig(1.5)- Indexing OLAP data using bitmap indices.

Logical operations such as AND and OR can be applied directly to bitmaps, drastically reducing I/O and computation time.

Below Figs-(4.6 &4.7)

Join indexing. In Example 3.4, we defined a star schema for *AllElectronics* of the form “*sales_star* [*time*, *item*, *branch*, *location*]: *dollars_sold* = sum (*sales_in_dollars*).” An example of a join index relationship between the *sales* fact table and the *location* and *item* dimension tables is shown in Figure 4.16. For example, the “Main Street” value in the *location* dimension table joins with tuples T57, T238, and T884 of the *sales* fact table. Similarly, the “Sony-TV” value in the *item* dimension table joins with tuples T57 and T459 of the *sales* fact table. The corresponding join index tables are shown in Figure 4.17.



Linkages between a *sales* fact table and *location* and *item* dimension tables.

Join index table for
location/sales

<i>location</i>	<i>sales_key</i>
...	...
Main Street	T57
Main Street	T238
Main Street	T884
...	...

Join index table for
item/sales

<i>item</i>	<i>sales_key</i>
...	...
Sony-TV	T57
Sony-TV	T459
...	...

Join index table linking
location and *item* to *sales*

<i>location</i>	<i>item</i>	<i>sales_key</i>
...
Main Street	Sony-TV	T57
...

Join index tables based on the linkages between the *sales* fact table and the *location* and *item* dimension tables shown in Figure 4.16.

Join Indexing

Join indexing captures relationships between fact tables and dimension tables. It records which fact table rows correspond to specific dimension values, avoiding costly joins during query execution.

Bitmap and join indexing can also be combined to form **bitmapped join indices**, further improving performance.

Processing OLAP Queries Efficiently

Efficient query processing involves selecting the most appropriate materialized cuboid, transforming query operations into cube operations, and exploiting available indices.

The system evaluates candidate cuboids based on granularity, selection predicates, and estimated cost, choosing the one that minimizes processing time.

OLAP Server Architectures

OLAP servers differ in how they store and manage data:

- **Relational OLAP systems** store data in relational tables and rely on SQL extensions and middleware for multidimensional processing.
- **Multidimensional OLAP systems** use array-based storage optimized for fast access to precomputed aggregates.
- **Hybrid OLAP systems** combine both approaches, storing detailed data relationally and aggregates multidimensionally.

Each architecture reflects a trade-off between scalability, performance, and storage efficiency.

Data Generalization and Concept Description

Data generalization summarizes data by replacing low-level values with higher-level concepts or by reducing dimensionality. This supports concise descriptions of large data sets and enables users to focus on general behavior rather than individual records.

Concept description is a form of data mining that produces **characterizations** and **comparisons** of data collections. It goes beyond enumeration to generate meaningful summaries.

Attribute-Oriented Induction

Attribute-oriented induction is a query-driven approach to concept description. It generalizes task-relevant data by examining attribute distinctness and applying either **attribute removal** or **attribute generalization**.

Attributes with too many distinct values and no meaningful generalization are removed. Attributes with defined hierarchies are generalized upward. Identical generalized tuples are merged, and aggregate measures such as counts or sums are accumulated.



Generalization is controlled using thresholds to avoid overgeneralization or undergeneralization. Users can adjust these thresholds dynamically, effectively performing roll-up and drill-down operations during mining.

The result is a compact, interpretable representation of the data that supports both descriptive analysis and deeper discovery.

Therefore, from star-schema-based querying to multidimensional mining and attribute-oriented induction, modern data warehousing provides a rich analytical ecosystem. Its strength lies not only in storing large volumes of data, but in enabling users to navigate, summarize, generalize, and discover knowledge interactively. When designed with business objectives in mind and implemented with efficient architectures, data warehouses form a critical foundation for advanced analytics and informed decision making.

OLAP Server Architectures: ROLAP, MOLAP, and HOLAP

OLAP servers present multidimensional views of data to users, abstracting away physical storage details. Internally, however, their architectures differ in how data is stored, aggregated, and accessed.



ROLAP (Relational OLAP)

- Uses **relational or extended-relational DBMS**.

- Stores data in **fact tables and summary fact tables**.
- OLAP operations are mapped to **SQL group-by queries**.
- Supports **base fact tables** (detailed data) and **summary fact tables** (aggregated data, using `all` for subtotals).
- **Highly scalable**, suitable for very large datasets.
- Slower aggregation compared to MOLAP.
- Common in enterprise-scale systems.

MOLAP (Multidimensional OLAP)

- Uses **array-based multidimensional storage**.
- Stores data directly in **data cube structures**.
- Provides **very fast query performance** due to precomputed aggregates.
- Suffers from **storage inefficiency for sparse data**.
- Uses **compression** and **two-level storage** (dense subcubes as arrays, sparse subcubes compressed).

HOLAP (Hybrid OLAP)

- Combines **ROLAP scalability** with **MOLAP speed**.
- Detailed data stored relationally; aggregates stored multidimensionally.
- Balances performance and storage efficiency.
- Common in commercial systems.

Specialized SQL Servers

- Optimized relational engines for OLAP-style queries.
- Support star/snowflake schemas in **read-only analytical environments**.

Data Storage in ROLAP vs MOLAP

- **ROLAP**: Data stored in relational tables; summaries represented by generalized attribute values (e.g., `day = all`).
- **MOLAP**: Data stored as multidimensional arrays.
- Warehouses follow **client-server architecture**; relational stores reside at the server, multidimensional stores may reside at server or client.

Data Generalization and Concept Description

Data Generalization

- Summarizes data by:
 - Replacing low-level values with higher-level concepts.
 - Reducing dimensionality.

- Enables analysis at **multiple abstraction levels**.

Concept Description

- A form of descriptive data mining.
- Produces:
 - **Characterization**: summarizes one target class.
 - **Comparison (Discrimination)**: contrasts a target class with others.

Limitations of Pure OLAP-Based Generalization

- Restricted mainly to **numeric measures** and **simple dimensions**.
- Weak support for **complex data types** (text, spatial, multimedia).
- Fully **user-driven**, requiring many manual OLAP operations.

Attribute-Oriented Induction (AOI)

AOI is a **query-driven, online generalization technique** that complements OLAP.

Core Ideas

- Collect **task-relevant data** via query.
- Generalize data using:
 - **Attribute removal**: remove attributes with many distinct values and no useful hierarchy.
 - **Attribute generalization**: climb concept hierarchies.
- Merge identical generalized tuples and accumulate:
 - count, sum, avg.

Generalization Control

- **Attribute threshold control**: limits distinct values per attribute.
- **Relation threshold control**: limits number of generalized tuples.
- Prevents overgeneralization and undergeneralization.

Key Properties

- Works on **complex data types**.
- No need for precomputed cubes.
- Supports **automated attribute relevance filtering**.
- Limited drill-down capability beyond generalized level.

Attribute-Oriented Induction for Class Comparison

- Compares a **target class** with **contrasting class(es)**.
- Uses **synchronous generalization** so all classes are compared at the same abstraction level.
- Produces quantitative contrasts (e.g., `count%`).
- Results are presented as tables, charts, or rules.

Points to remember:

- **ROLAP**: scalable, SQL-based, slower aggregation.
- **MOLAP**: fast, cube-based, storage-heavy.
- **HOLAP**: combines strengths of both.
- **OLAP** excels at interactive summarization.
- **Attribute-oriented induction** enables automated, flexible, online generalization.
- **Best practice**: integrate OLAP and AOI for balanced performance, flexibility, and depth of analysis.

Dimensions and Concept Hierarchies

Dimensions often include hierarchical relationships that allow data to be analyzed at varying levels of abstraction. For example, locations may roll up from city to state to country, and time may roll up from day to month to year.

Hierarchies can be total orders, where each level maps neatly to a higher level, or partial orders, where multiple aggregation paths exist. They may also be defined by grouping values into ranges, such as price bands. These hierarchies enable flexible analysis through operations that summarize or drill into data.

Measures and Their Computation

Measures in a data cube are computed using aggregation functions and can be classified based on how they behave under aggregation.

- **Distributive measures** can be computed incrementally from subaggregates, such as sum, count, minimum, and maximum.
- **Algebraic measures** are derived from a fixed number of distributive measures, such as average or standard deviation.

- **Holistic measures** require access to the full data set or a large portion of it, such as median or mode, and are more difficult to compute efficiently.

Understanding these categories is essential for designing efficient aggregation strategies and scalable analytical systems.

A data warehouse is both a repository and an architecture—a foundation for analytical reasoning across an enterprise. By separating analytical workloads from operational systems, integrating heterogeneous data sources, and organizing information in multidimensional structures, data warehousing enables organizations to transform raw data into actionable insight. As analytical needs grow and technologies evolve, the principles underlying data warehousing continue to remain central to effective decision support.

15th January, 2026