What are Operational DBMS?

- They consist of Tables having attributes and are populated by tuples.
- They generally use the E-R data model.
- It is used to store transactional data.
- The information content is generally recent.
- These are thus called as OLTP systems.
- Their goals are data accuracy & consistency , Concurrency , Recoverability, Reliability (ACID Properties).

Inmon's Definition of a Data Warehouse

A **data warehouse** is a

- subject-oriented,
- integrated,
- nonvolatile, and
- time-variant

collection of data in support of management's decisions. The data warehouse contains granular corporate data.

Data Warehouse Properties

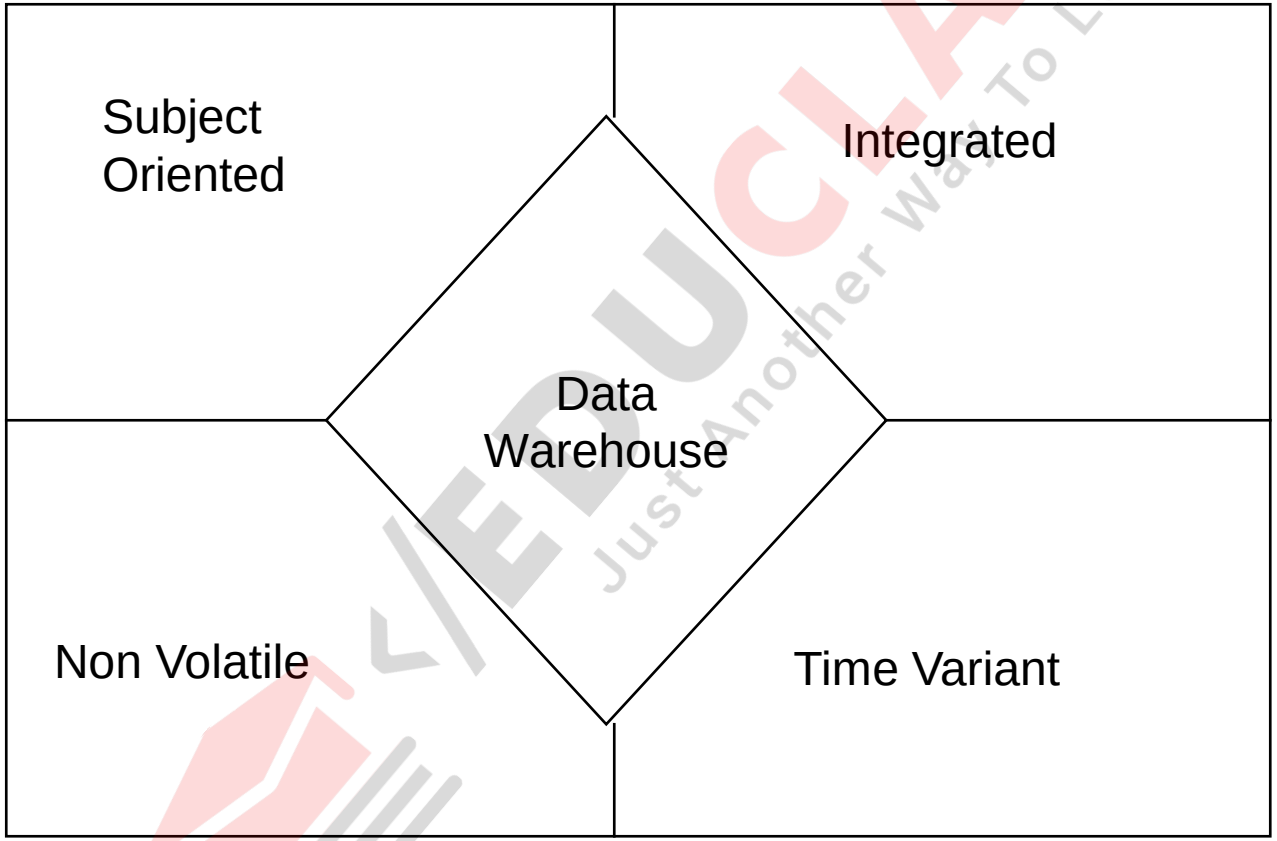

Subject-Oriented

Data is categorized and stored by business subject rather than by application

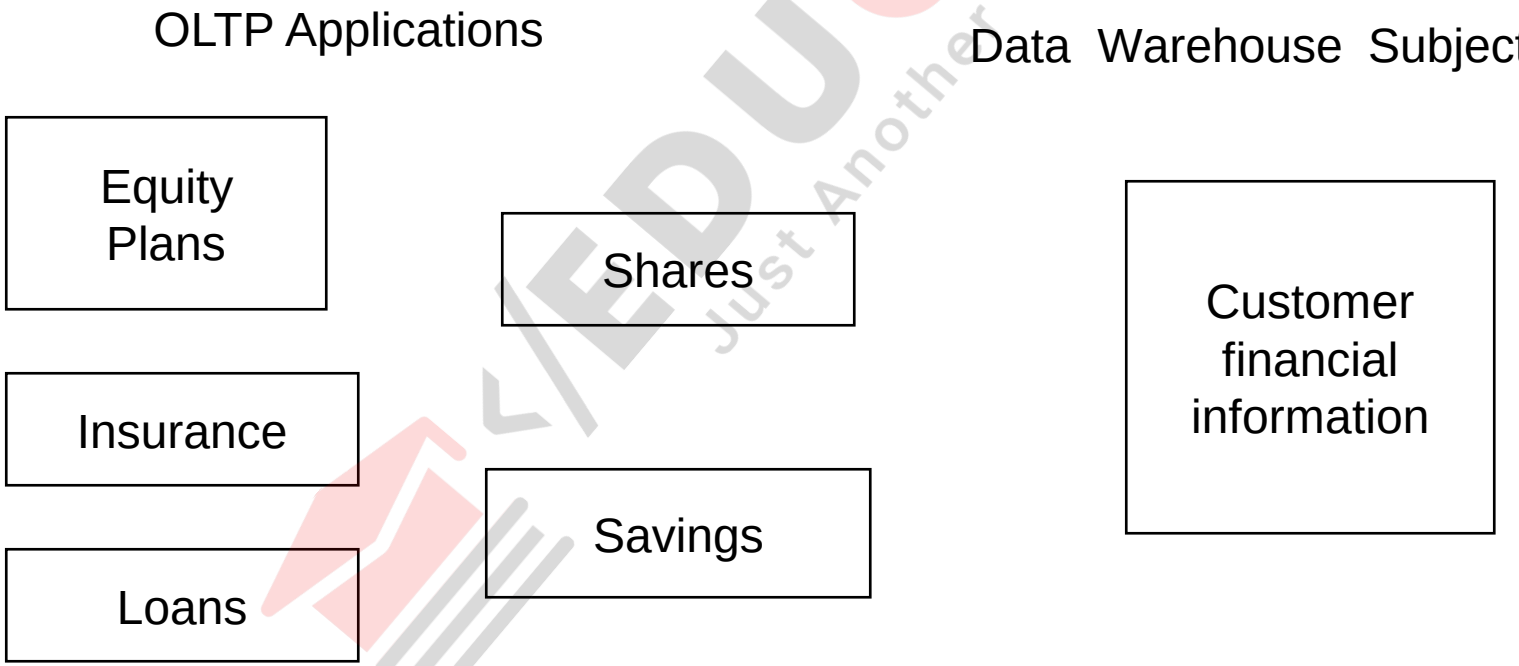

Subject-Oriented Data Collections

- Organized around major subjects, such as customer, product, sales.

- Focusing on the modeling and analysis of data for decision makers, not on daily operations or transaction processing.

- Provide a simple and concise view around particular subject issues by excluding data that are not useful in the decision support process.

Subject-Oriented Data Collections

- For a manufacturer, the major subject areas might be product, order, vendor, bill of material, and raw goods. For a retailer, the major subject areas may be product, sale, vendor, and so forth. Each type of company has its own unique set of subjects.

Subject Oriented

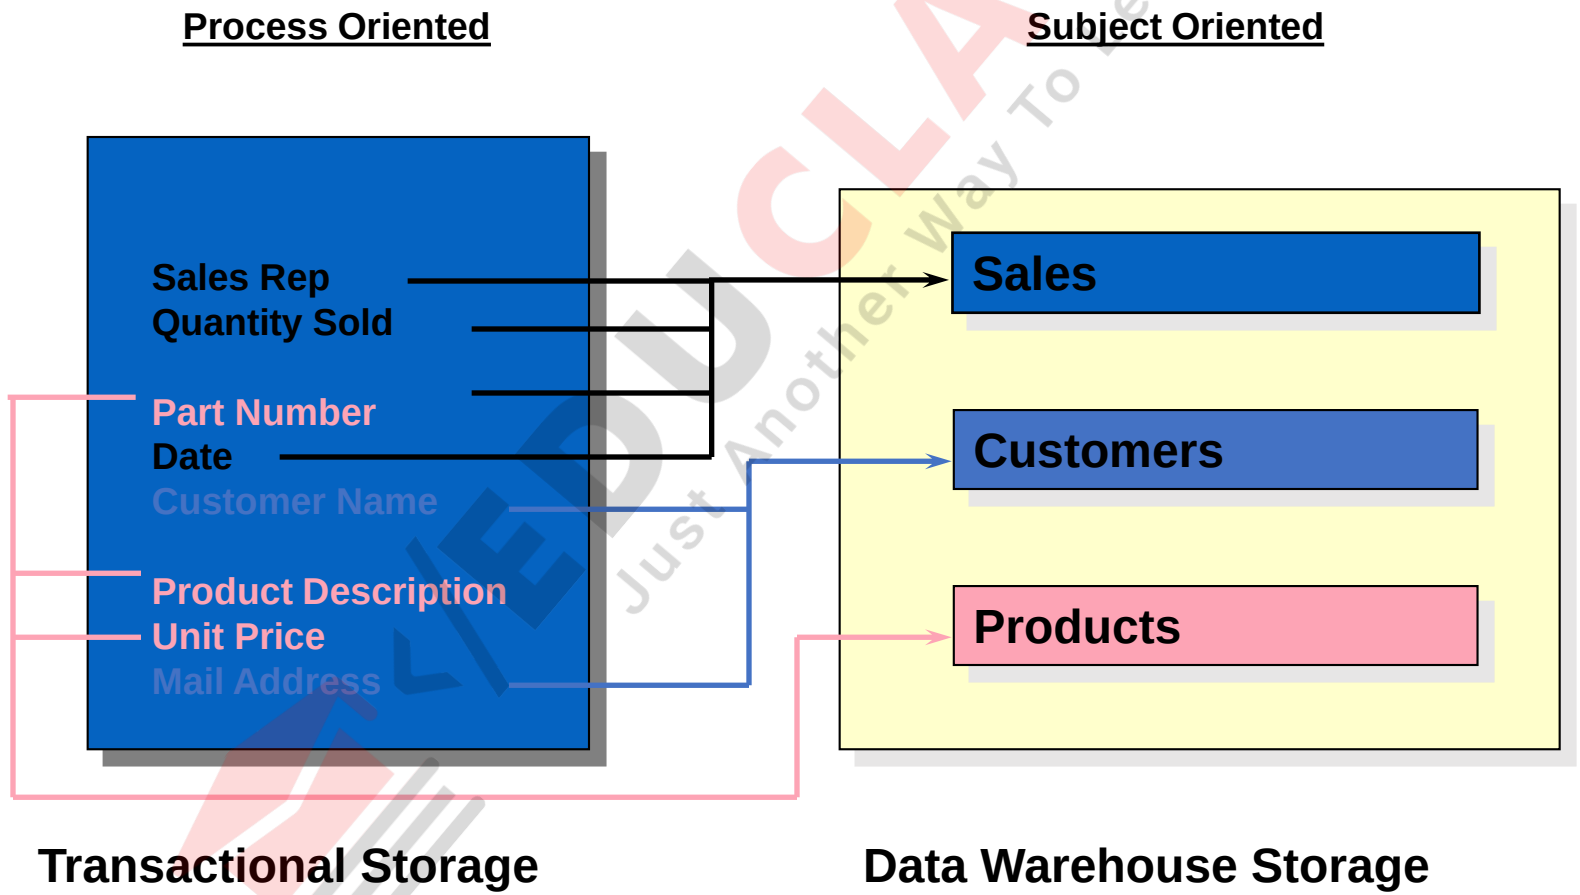

Integrated

Data on a given subject is defined and stored once.

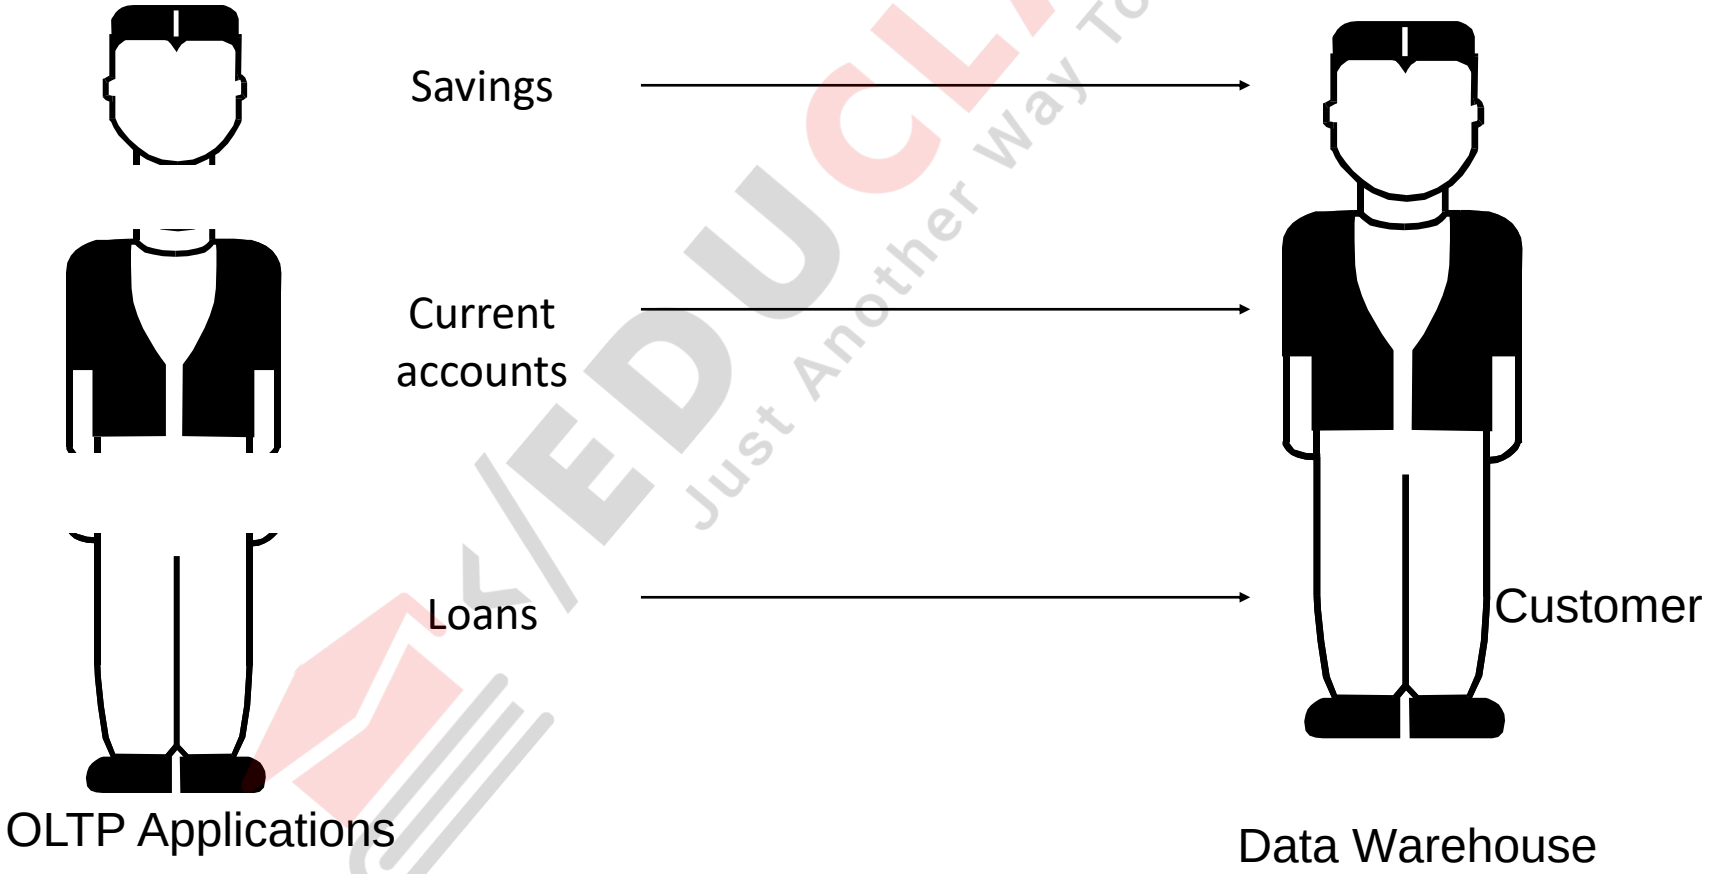

Integrated Data Collections

- Of all the aspects of a data warehouse, integration is the most important. Data is fed from multiple disparate sources into the data warehouse.
- As the data is fed it is converted, reformatted, resequenced, summarized, and so forth. The result is that data—once it resides in the data warehouse—has a single physical corporate image.

Integration of Data

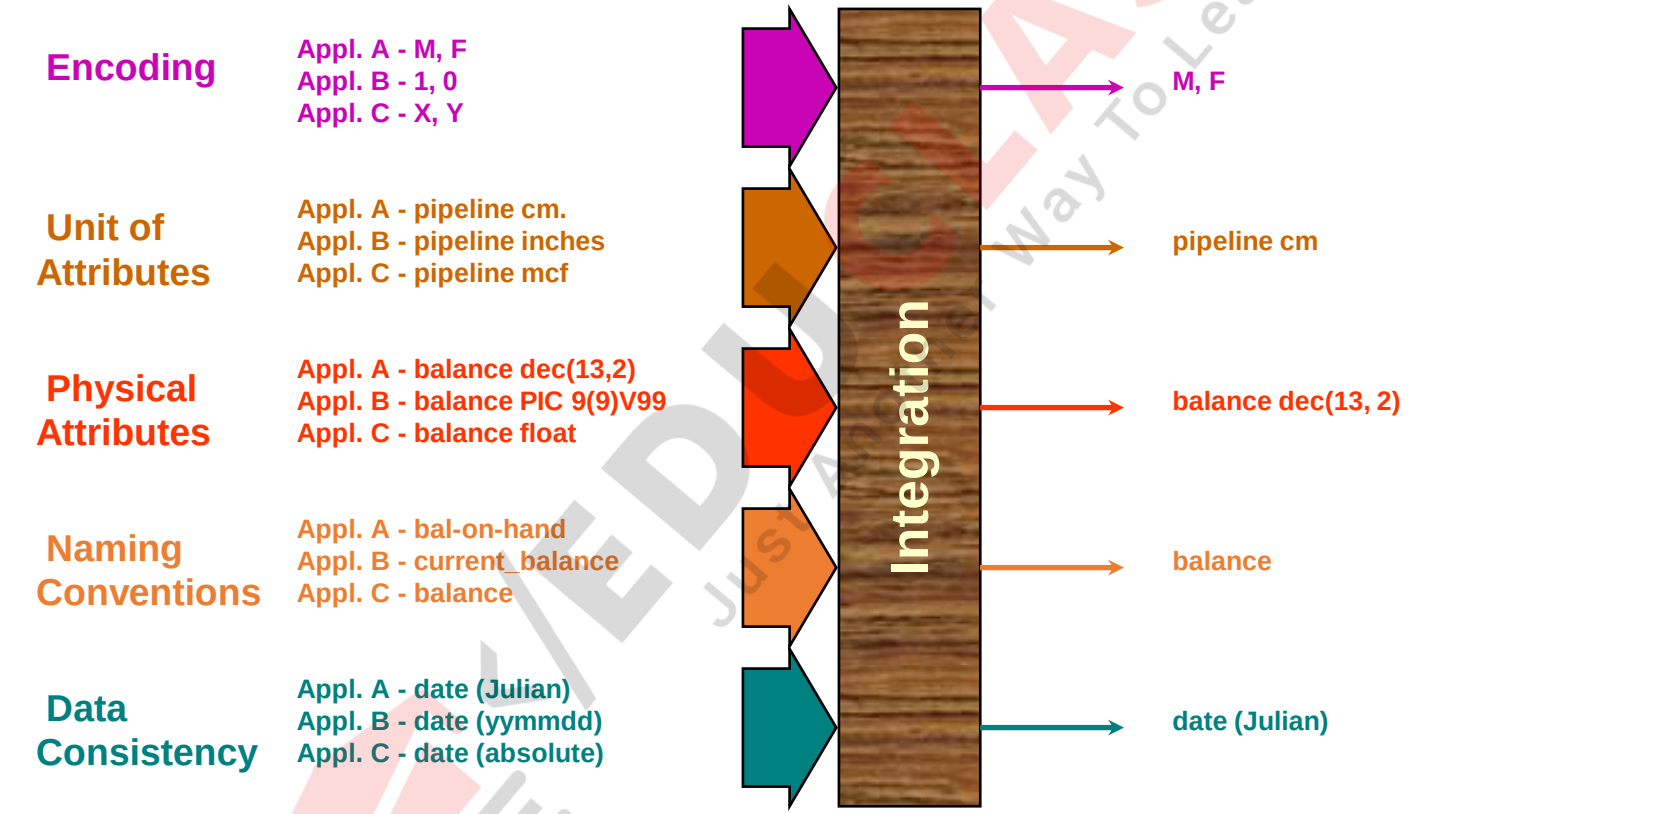

Volatility of Data

Volatile	Non-Volatile
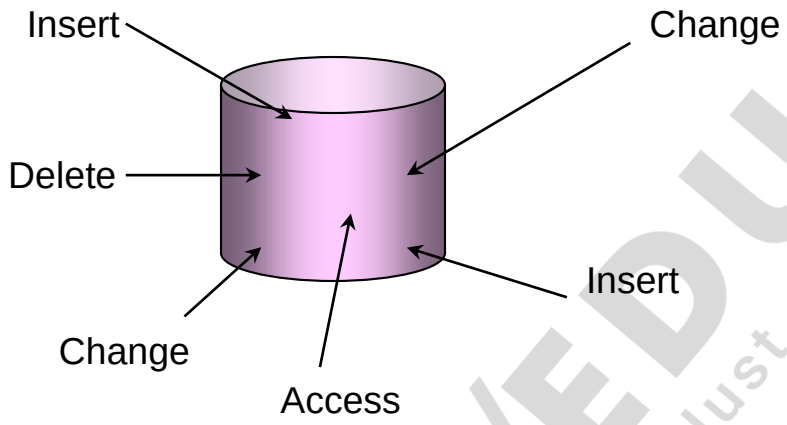	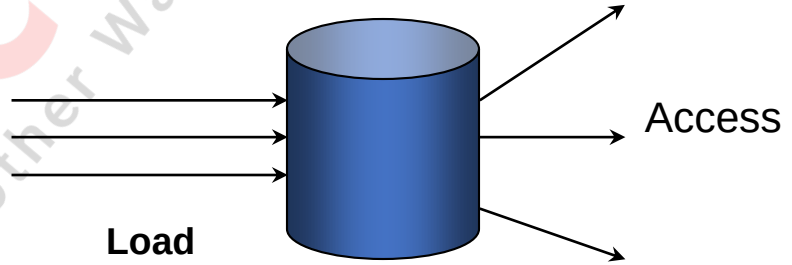
Record-by-Record Data Manipulation	Mass Load / Access of Data
Transactional Storage	**Data Warehouse Storage**

Nonvolatile

Typically data in the data warehouse is not updated or deleted.

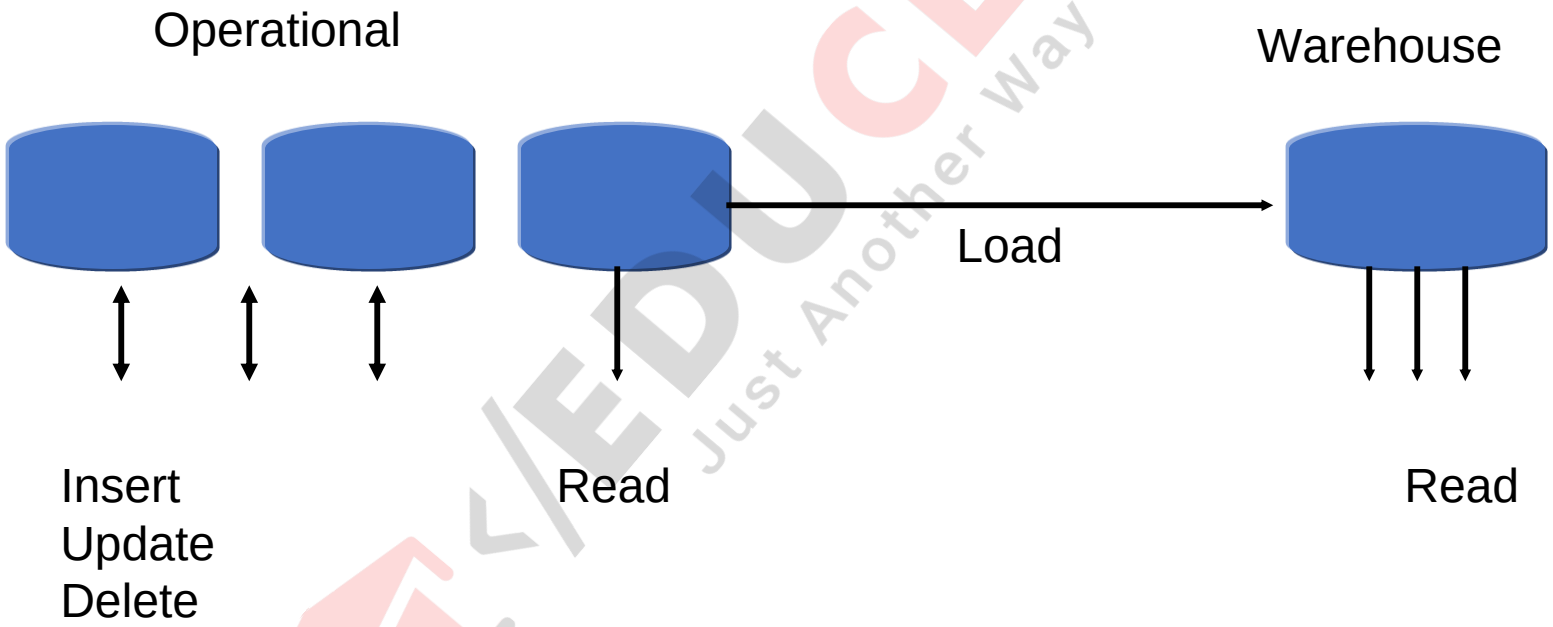

Non-volatile Data Collections

Data is updated in the operational environment as a regular matter of course, but warehouse data exhibits a very different set of characteristics. Data warehouse data is loaded (usually en masse) and accessed, but it is not updated (in the general sense).

Non-volatile Data Collections

Instead, when data in the data warehouse is loaded, it is loaded in a snapshot, static format. When subsequent changes occur, a new snapshot record is written. In doing so a history of data is kept in the data warehouse.

Time-variant Data Collections

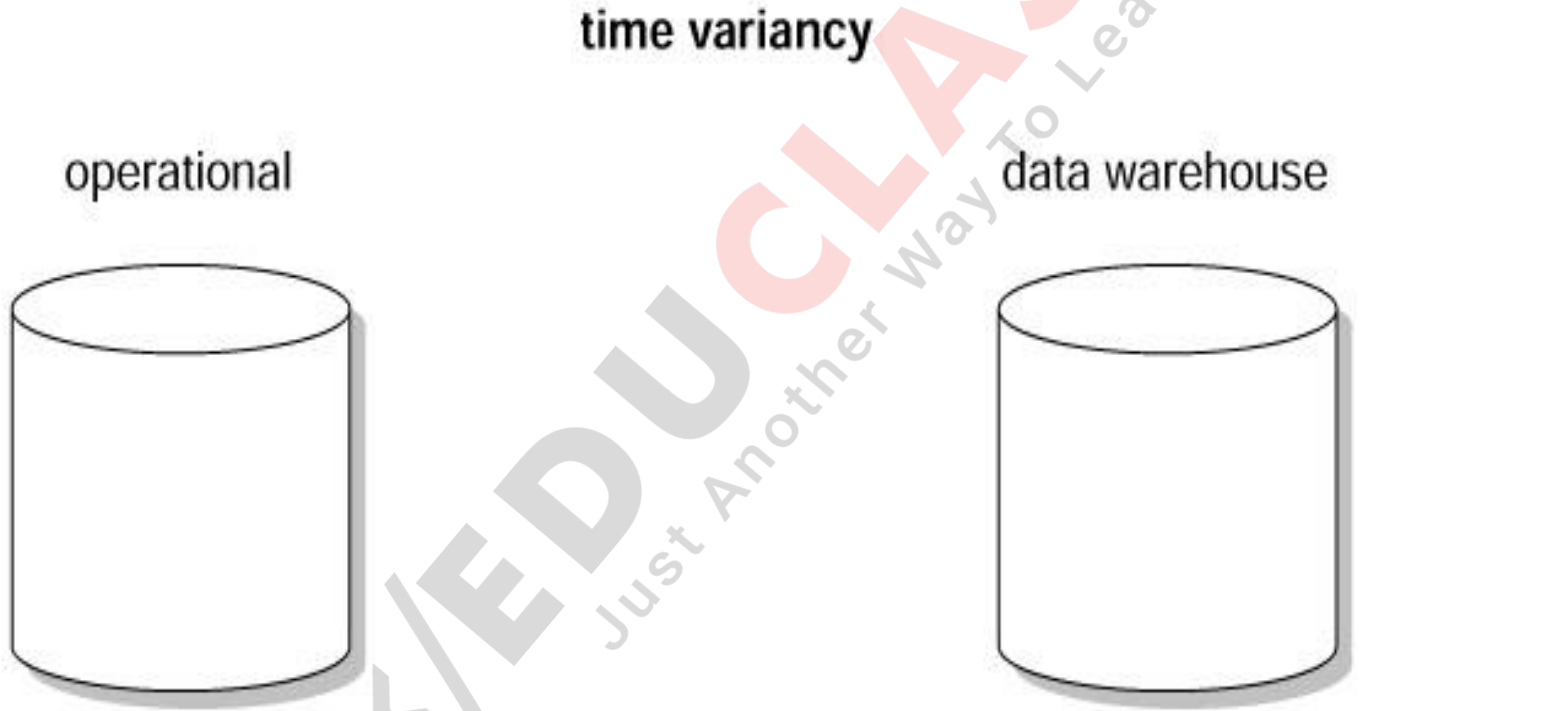

- time horizon—current to 60–90 days
- update of records
- key structure may/may not contain an element of time

- time horizon—5–10 years
- sophisticated snapshots of data
- key structure contains an element of time

Time Variant

Current Data

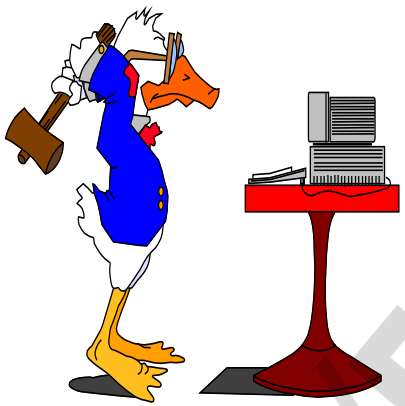

Historical Data

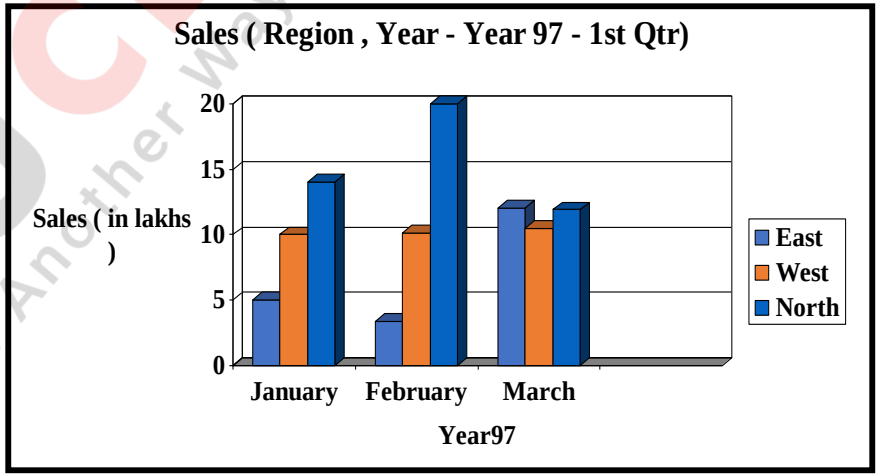

Transactional Storage

Data Warehouse Storage

Time-variant Data Collections

Time variance implies that every unit of data in the data warehouse is accurate as of some one moment in time. In some cases, a record is time stamped. In other cases, a record has a date of transaction. But in every case, there is some form of time marking to show the moment in time during which the record is accurate.

Time-variant Data Collections

A 60-to-90-day time horizon is normal for operational systems; a 5-to-10-year time horizon is normal for the data warehouse. As a result of this difference in time horizons, the data warehouse contains *much* more history than any other environment.

What are Data Warehouses?

- Data warehouses store large volumes of data which are frequently used by DSS
- It is maintained separately from the organization's operational databases
- Data warehouses are relatively static with only infrequent updates
- A data warehouse is a stand-alone repository of information, integrated from several, possibly heterogeneous operational databases

Data Warehouse Architecture

- There are 2 approaches for constructing data-warehouse:
  - Top-down approach and
  - Bottom-up approach

Top-down approach

- This approach is defined by **Inmon** as – datawarehouse as a central repository for the complete organisation and data marts are created from it after the complete datawarehouse has been created.

Top-down approach

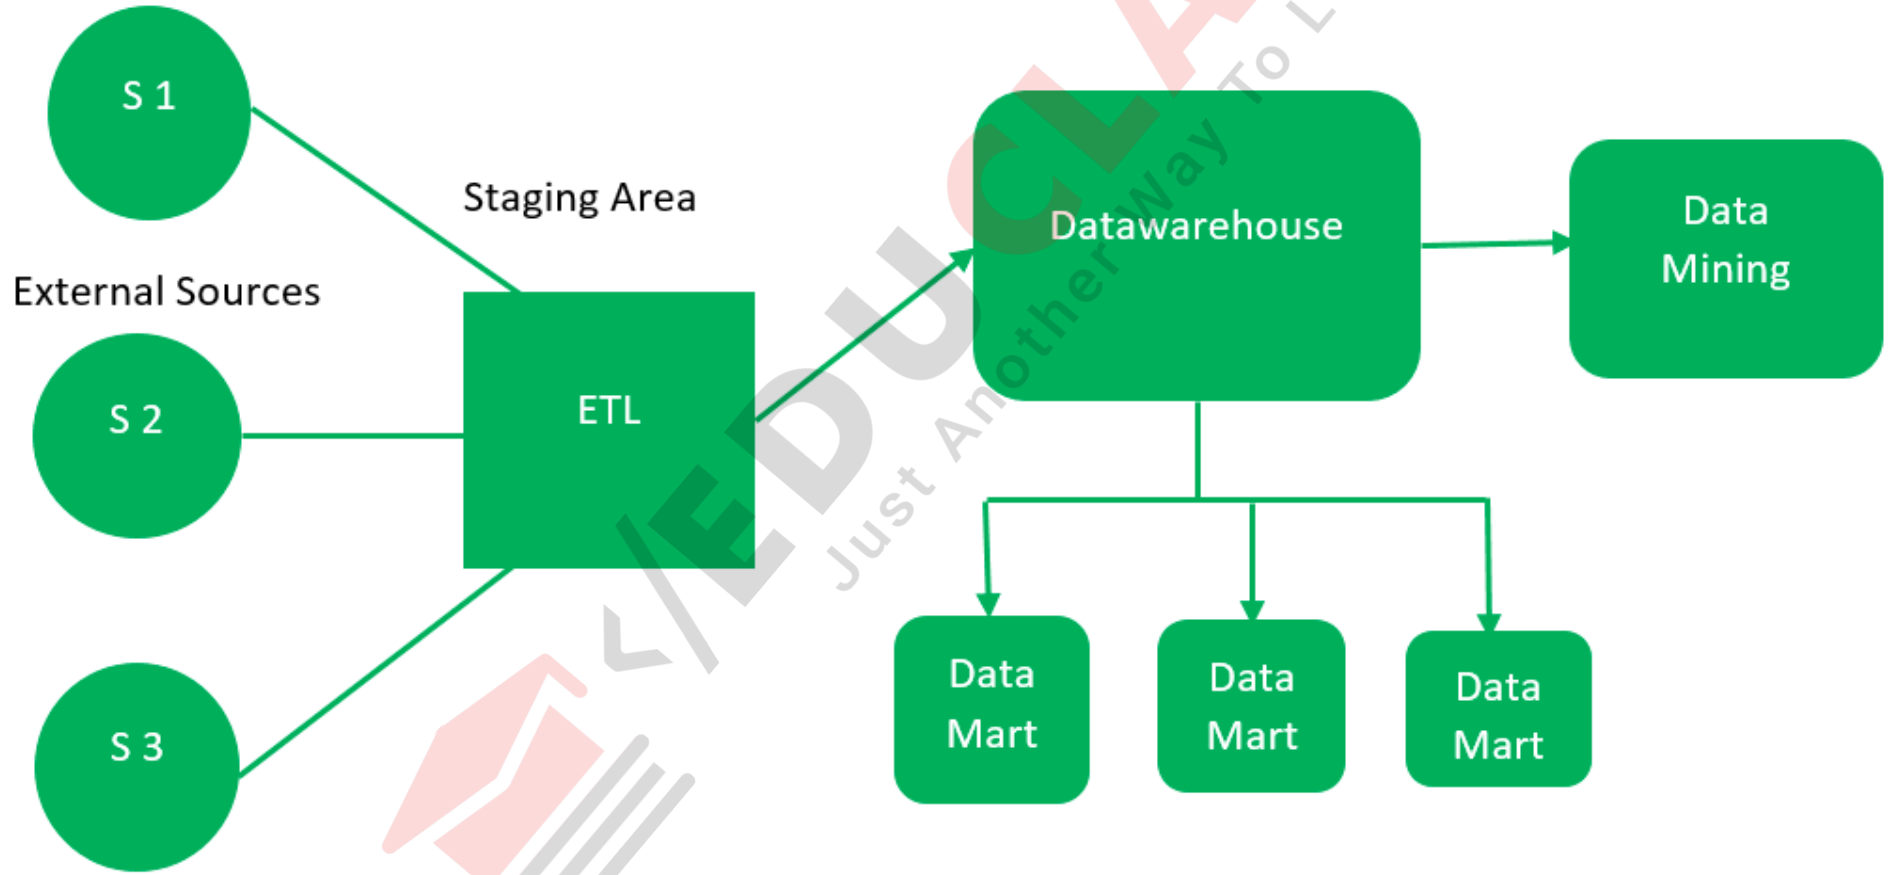

- **Advantages of Top-Down Approach –**
- Since the data marts are created from the data warehouse, provides consistent dimensional view of data marts.
- Also, this model is considered as the strongest model for business changes. That’s why, big organizations prefer to follow this approach.
- Creating data mart from data warehouse is easy.
- Centralized rules and control

- **Disadvantages of Top-Down Approach –**
- The cost, time taken in designing and its maintenance is very high.
- High exposure to risk of failure.
- Needs high level of cross functional skills.

Bottom-up approach

- This approach is given by **Kinball** as – data marts are created first and provides a thin view for analyses and datawarehouse is created after complete data marts have been created.

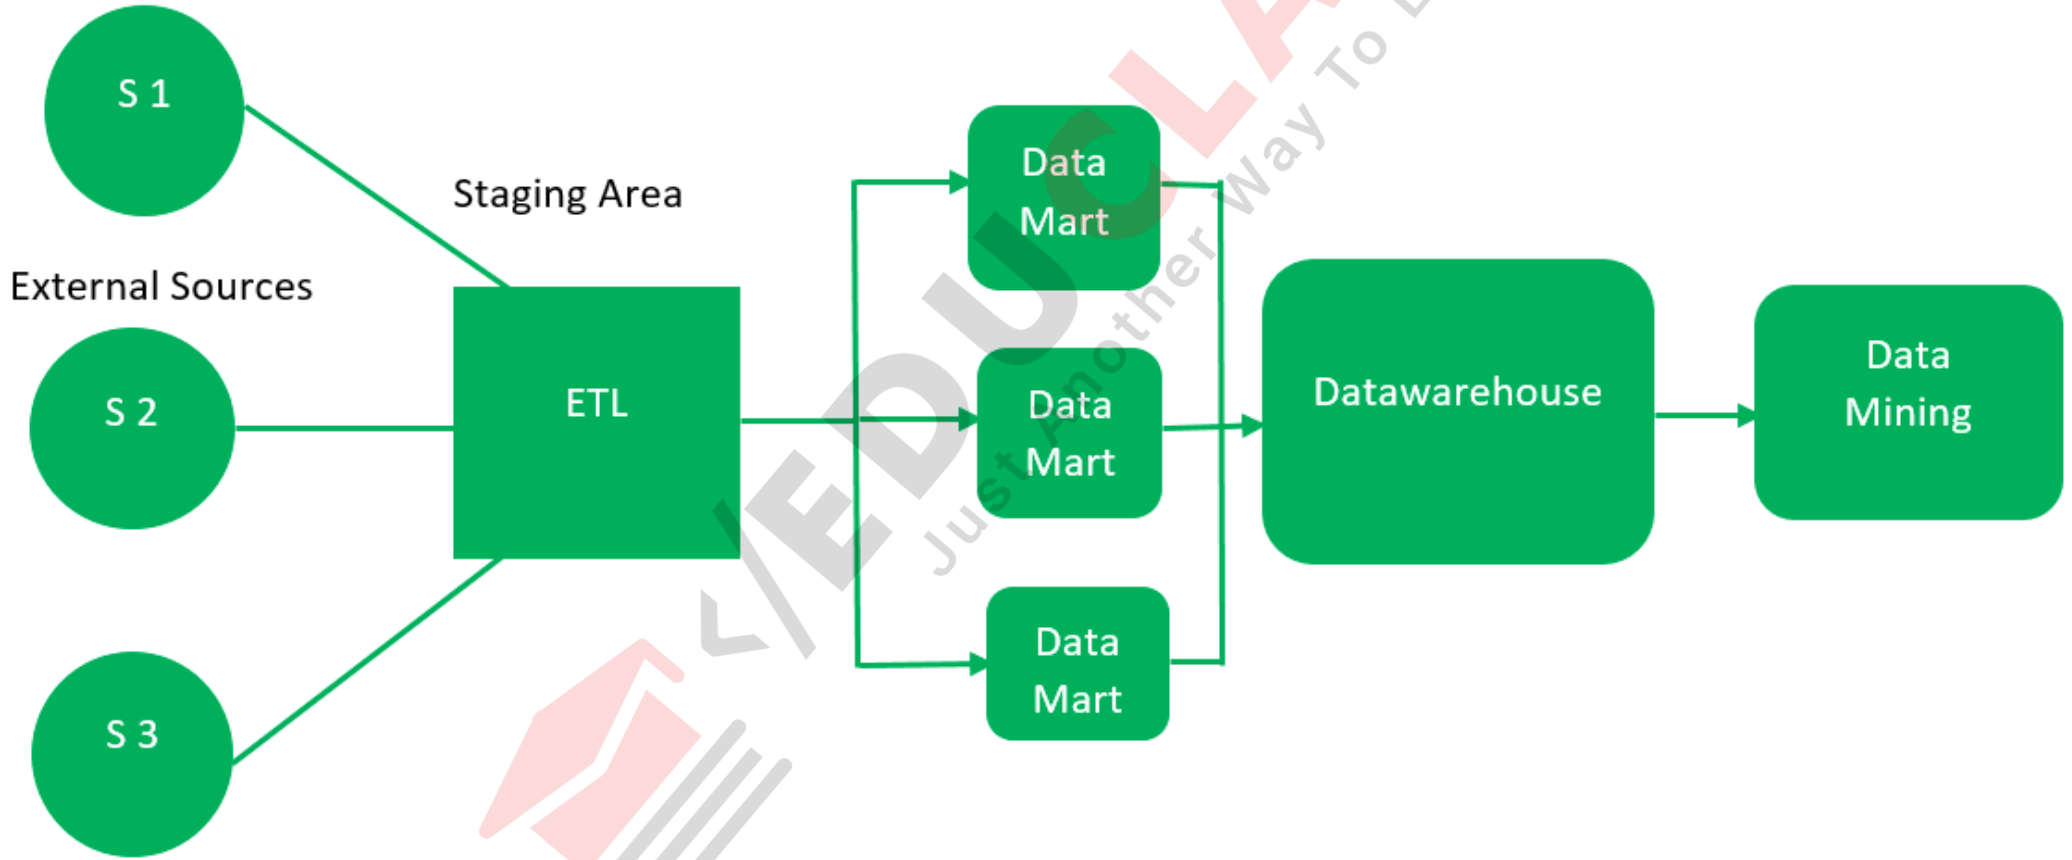

- First, the data is extracted from external soures (same as happens in top-down approach).
- Then, the data go through the staging area (as explained above) and loaded into data marts instead of datawarehouse. The data marts are created first and provide reporting capability. It addresses a single business area.
- These data marts are then integrated into datawarehouse.

- **Advantages of Bottom-Up Approach –**
- As the data marts are created first, so the reports are quickly generated.
- We can accomodate more number of data marts here and in this way datawarehouse can be extended.
- Also, the cost and time taken in designing this model is low comparatively.
- Less risk of failure.
- Allows project team to learn and grow.

- **Disadvantage of Bottom-Up Approach –**
- This model is not strong as top-down approach as dimensional view of data marts is not consistent as it is in above approach.
- Permeates redundant data in every data mart.
- Perpetuates inconsistent and irreconcilable data.
- Proliferates unmanageable interfaces.

Architecture of DW

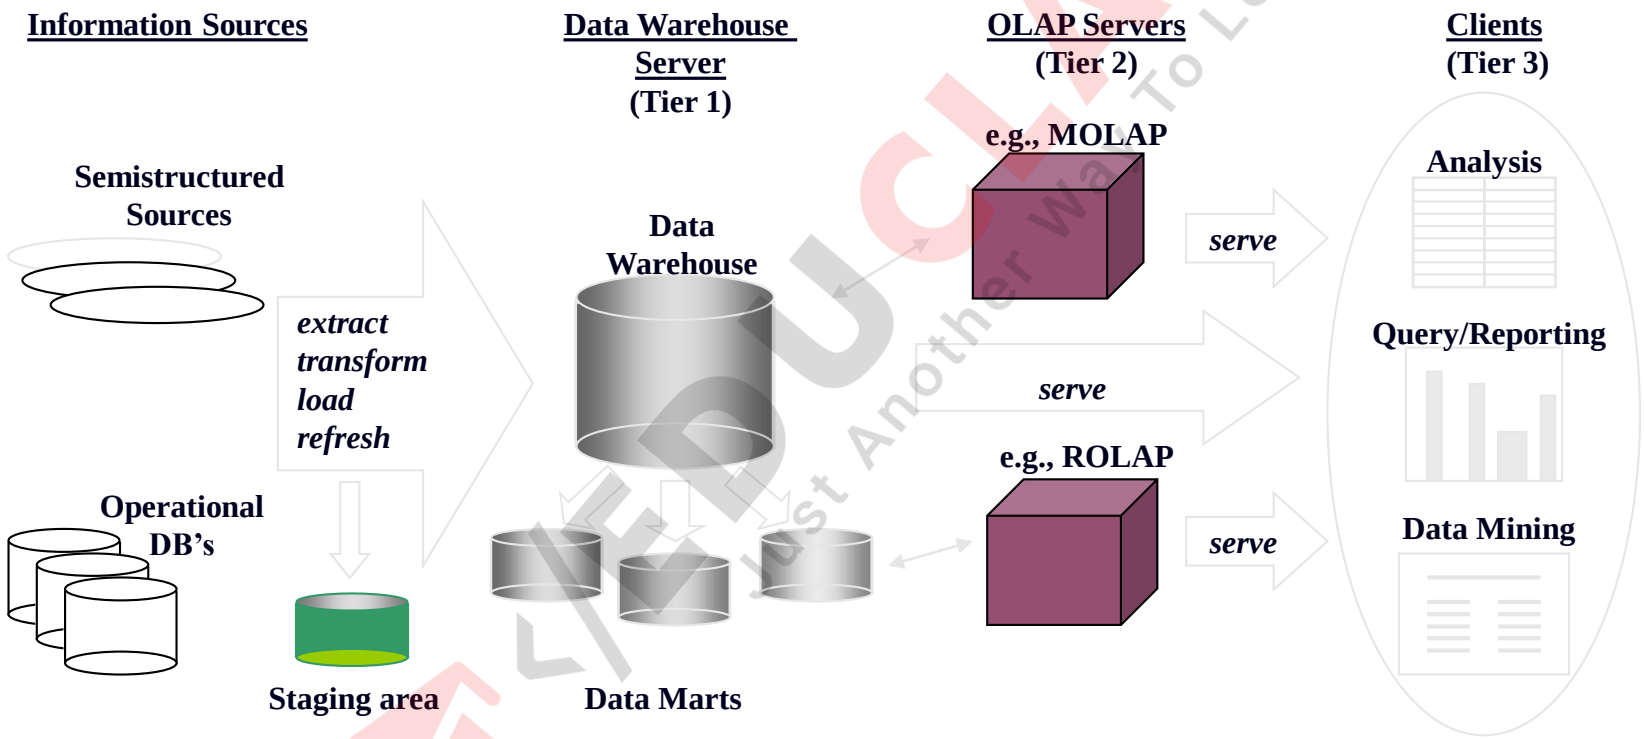

Components

- Major components
  - Source data component
  - Data staging component
  - Information delivery component
  - Metadata component
  - Management and control component

1. Source Data Components

- Source data can be grouped into 4 components
  - Production data
    - Comes from operational systems of enterprise
    - Some segments are selected from it
    - Narrow scope, e.g. order details
  - Internal data
    - Private datasheet, documents, customer profiles etc.
    - E.g. Customer profiles for specific offering
    - Special strategies to transform ‘it’ to DW (text document)
  - Archived data
    - Old data is archived
    - DW have snapshots of historical data
  - External data
    - Executives depend upon external sources
    - E.g. market data of competitors, car rental require new manufacturing. Define conversion

Architecture of DW

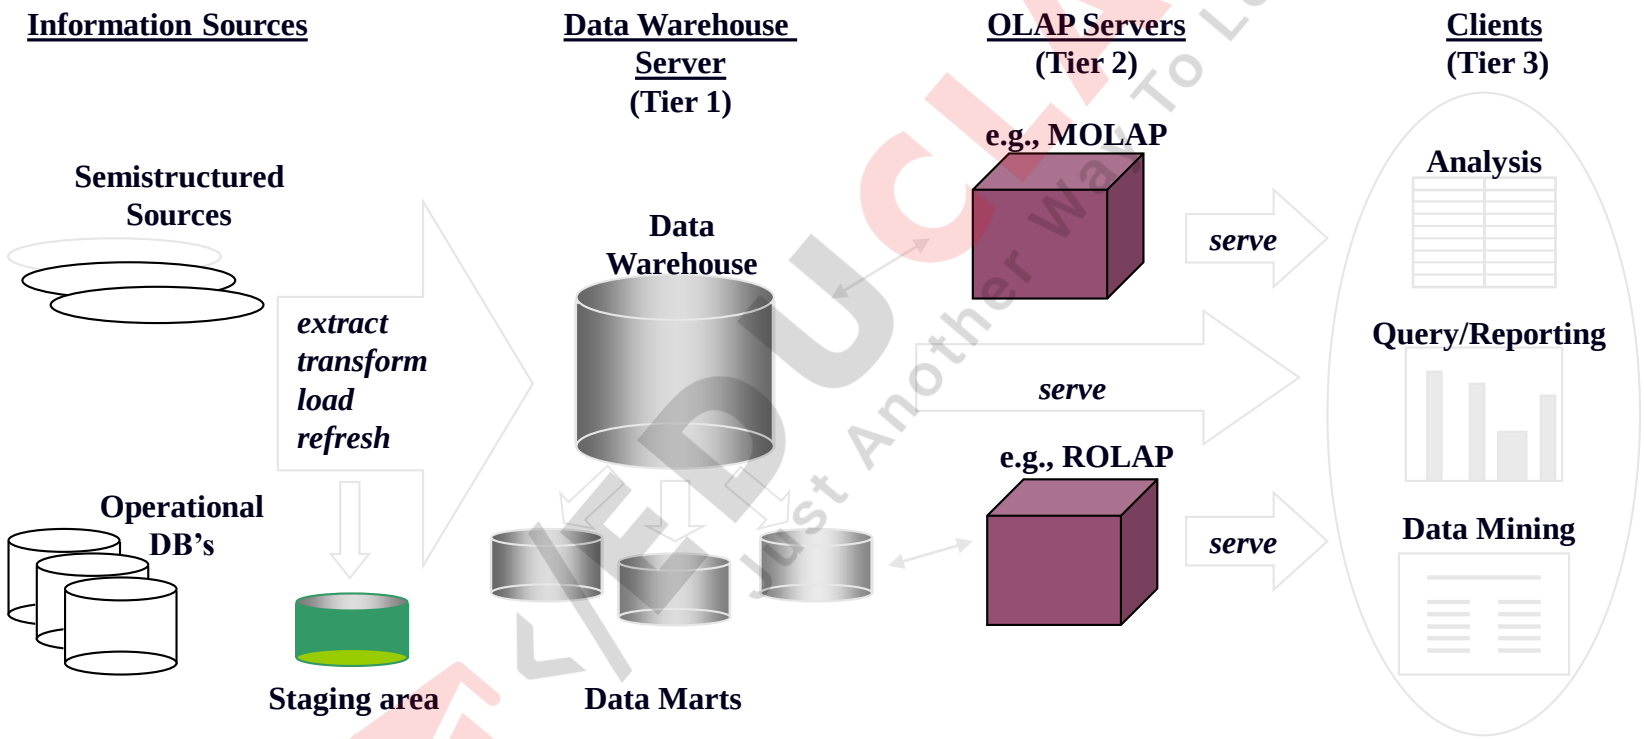

2. Data Staging Components

- After data is extracted, data is to be prepared
- Data extracted from sources needs to be changed, converted and made ready in suitable format
- Three major functions to make data ready
  - Extract
  - Transform
  - Load
- Staging area provides a place and area with a set of functions to
  - Clean
  - Change
  - Combine
  - Convert

Architecture of DW

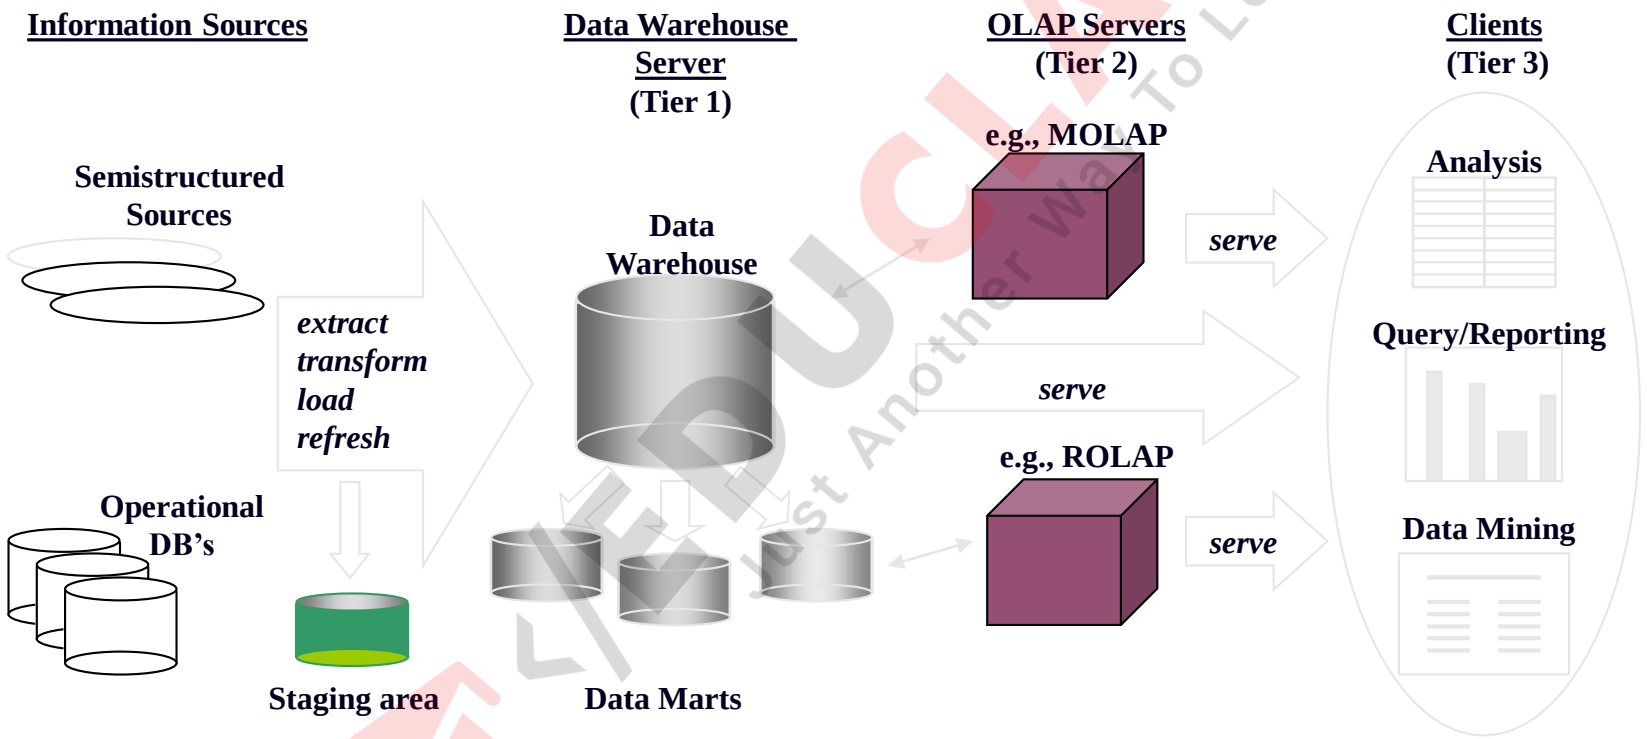

3. Data Storage Components

- Separate repository
- Data structured for efficient processing
- Redundancy is increased
- Updated after specific periods
- Only read-only

Architecture of DW

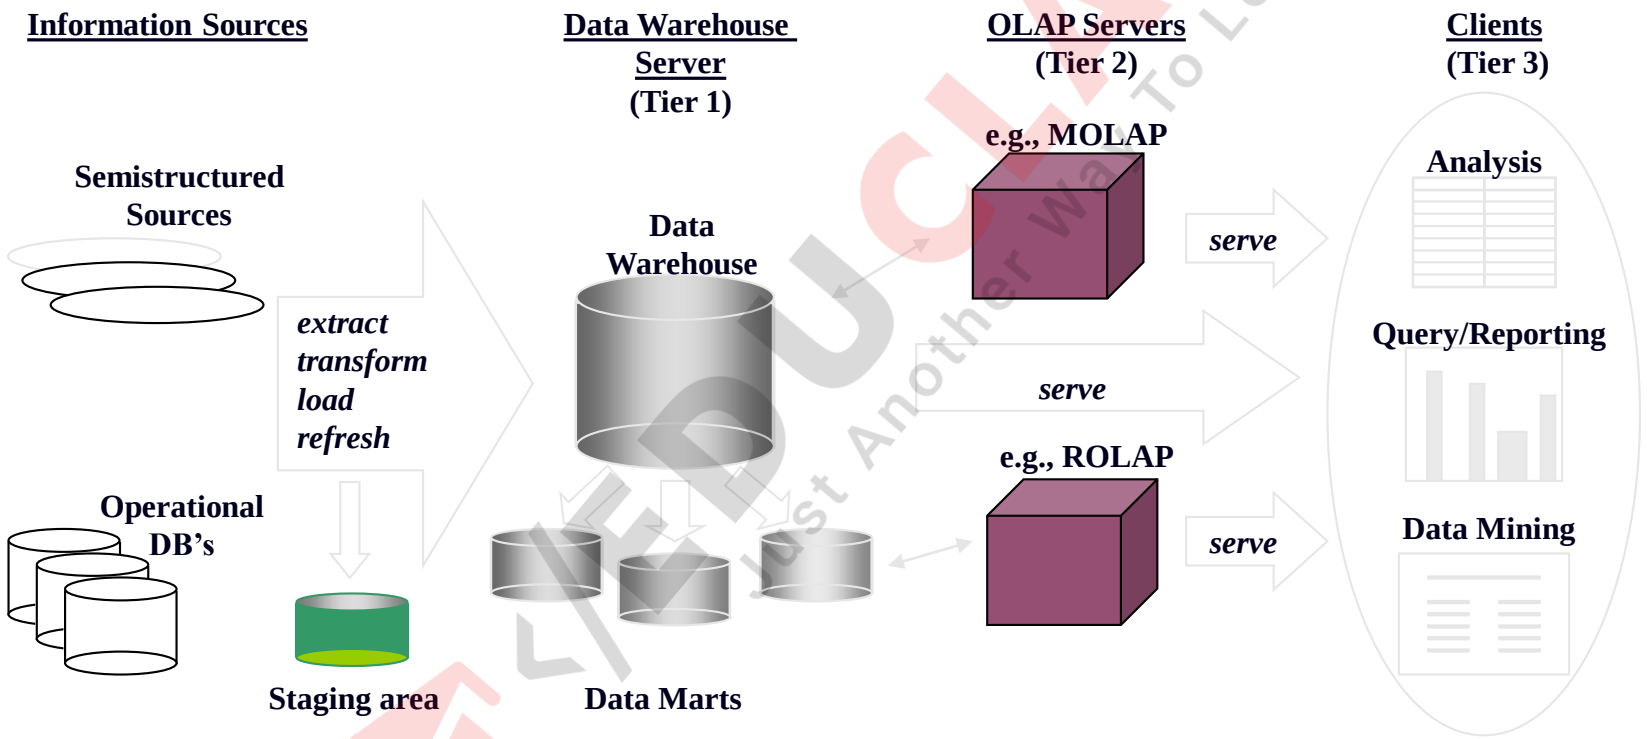

4. Information Delivery Component

- To provide information to the wide community of data warehouse users, the information delivery component includes different methods of information delivery.
- Ad hoc reports are predefined reports primarily meant for novice and casual users.
- Provision for complex queries, multidimensional (MD) analysis, and statistical analysis cater to the needs of the business analysts and power users.
- Information fed into Executive Information Systems (EIS) is meant for senior executives and high-level managers.
- Some data warehouses also provide data to data-mining applications. Data-mining applications are knowledge discovery systems where the mining algorithms help you discover trends and patterns from the usage of your data.

Metadata Component

- Metadata in a data warehouse is similar to the data dictionary or the data catalog in a database management system.
- In the data dictionary, you keep the information about the logical data structures, the information about the files and addresses, the information about the indexes, and so on.
- The data dictionary contains data about the data in the database.
- Similarly, the metadata component is the data about the data in the data warehouse.
- Metadata in a data warehouse is similar to a data dictionary, but much more than a data dictionary.

Management and Control Component

- This component of the data warehouse architecture sits on top of all the other components.
- The management and control component coordinates the services and activities within the data warehouse.
- This component controls the data transformation and the data transfer into the data warehouse storage.
- On the other hand, it moderates the information delivery to the users.
- It works with the database management systems and enables data to be properly stored in the repositories.
- It monitors the movement of data into the staging area and from there into the data warehouse storage itself.
- The management and control component interacts with the metadata component to perform the management and control functions.
- As the metadata component contains information about the data warehouse itself, the metadata is the source of information for the management module.

Data Mart

- A **Data Mart** is focused on a single functional area of an organization and contains a subset of data stored in a Data Warehouse.
- A Data Mart is a condensed version of Data Warehouse and is designed for use by a specific department, unit or set of users in an organization. E.g., Marketing, Sales, HR or finance. It is often controlled by a single department in an organization.
- Data Mart usually draws data from only a few sources compared to a data warehouse.
- Data marts are small in size and are more flexible compared to a Datawarehouse.

Types of Data Mart

- There are three main types of data mart:
- **Dependent**: Dependent data marts are created by drawing data directly from operational, external or both sources.
- **Independent**: Independent data mart is created without the use of a central data warehouse.
- **Hybrid**: This type of data marts can take data from data warehouses or operational systems.

Dependent Data Mart

- In a dependent data mart, data is sourced from the existing data warehouse itself. This is a top-down approach because the portion of restructured data into the data mart is extracted from the centralized data warehouse.
- **A data mart can use DW data either logically or physically as shown below:**
- **Logical View:** In this scenario, data mart's data is not physically separated from the DW. It refers to DW data through virtual views (or) tables logically.
- **Physical subset:** In this scenario, data mart's data is physically separated from the DW.
- Once one or more data marts are developed, you can allow the users to access only the data marts (or) to access both Data marts and Data warehouses.

Cont..

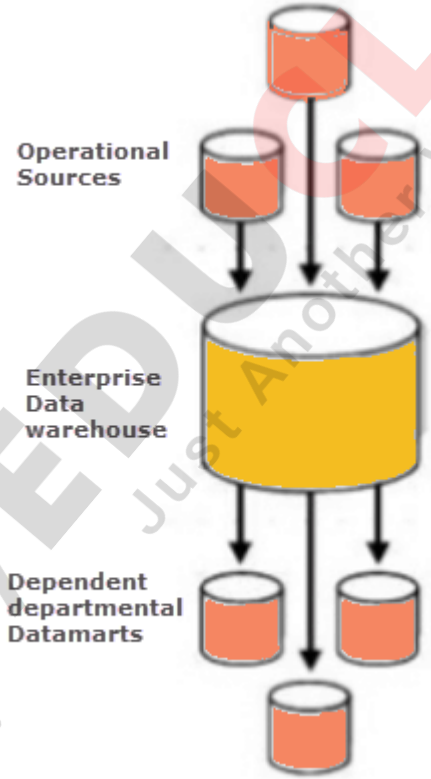

Independent Data Mart

- An independent data mart is best suitable for small departments in an organization. Here data is not sourced from the existing data warehouse. The Independent data mart is neither dependent on enterprise DW nor other data marts.
- Independent data marts are stand-alone systems where data is extracted, transformed and loaded from external (or) internal data sources. These are easy to design and maintain until it is supporting simple department wise business needs.

Cont..

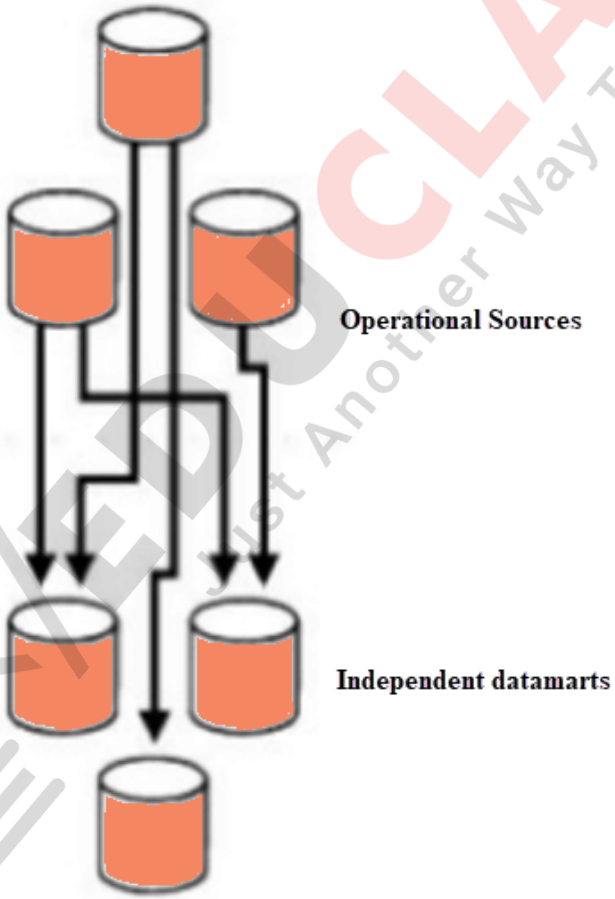

Hybrid Data Mart

- A hybrid data mart combines input from sources apart from Data warehouse. This could be helpful when you want ad-hoc integration, like after a new group or product is added to the organization.
- It is the best data mart example suited for multiple database environments and fast implementation turnaround for any organization. It also requires least data cleansing effort. Hybrid Data mart also supports large storage structures, and it is best suited for flexible for smaller data-centric applications.

Cont…

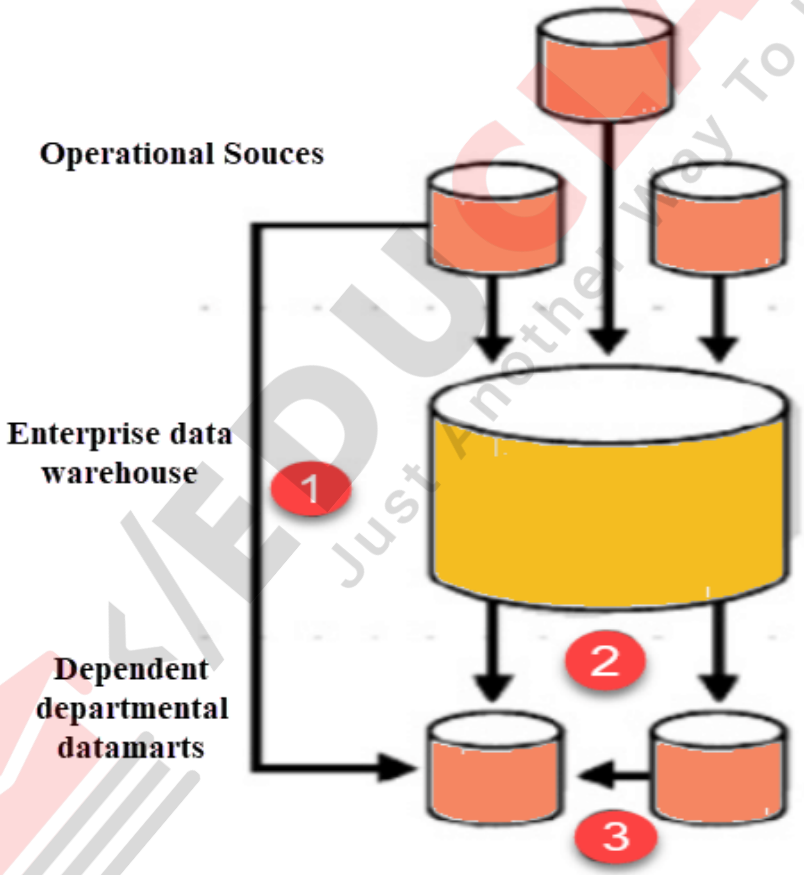

Advantages of Data Marts

- Data marts contain a subset of organization-wide data. This Data is valuable to a specific group of people in an organization.
- It is cost-effective alternatives to a data warehouse, which can take high costs to build.
- Data Mart allows faster access of Data.
- Data Mart is easy to use as it is specifically designed for the needs of its users. Thus a data mart can accelerate business processes.
- Data Marts needs less implementation time compare to Data Warehouse systems. It is faster to implement Data Mart as you only need to concentrate the only subset of the data.
- It contains historical data which enables the analyst to determine data trends.

Disadvantages of DataMart

- Many a times enterprises create too many disparate and unrelated data marts without much benefit. It can become a big hurdle to maintain.
- Data Mart cannot provide company-wide data analysis as their data set is limited.

What is ETL?

- ETL (Extraction, Transformation and Loading) is a process by which data is integrated and transformed from the operational systems into the Data Warehouse environment

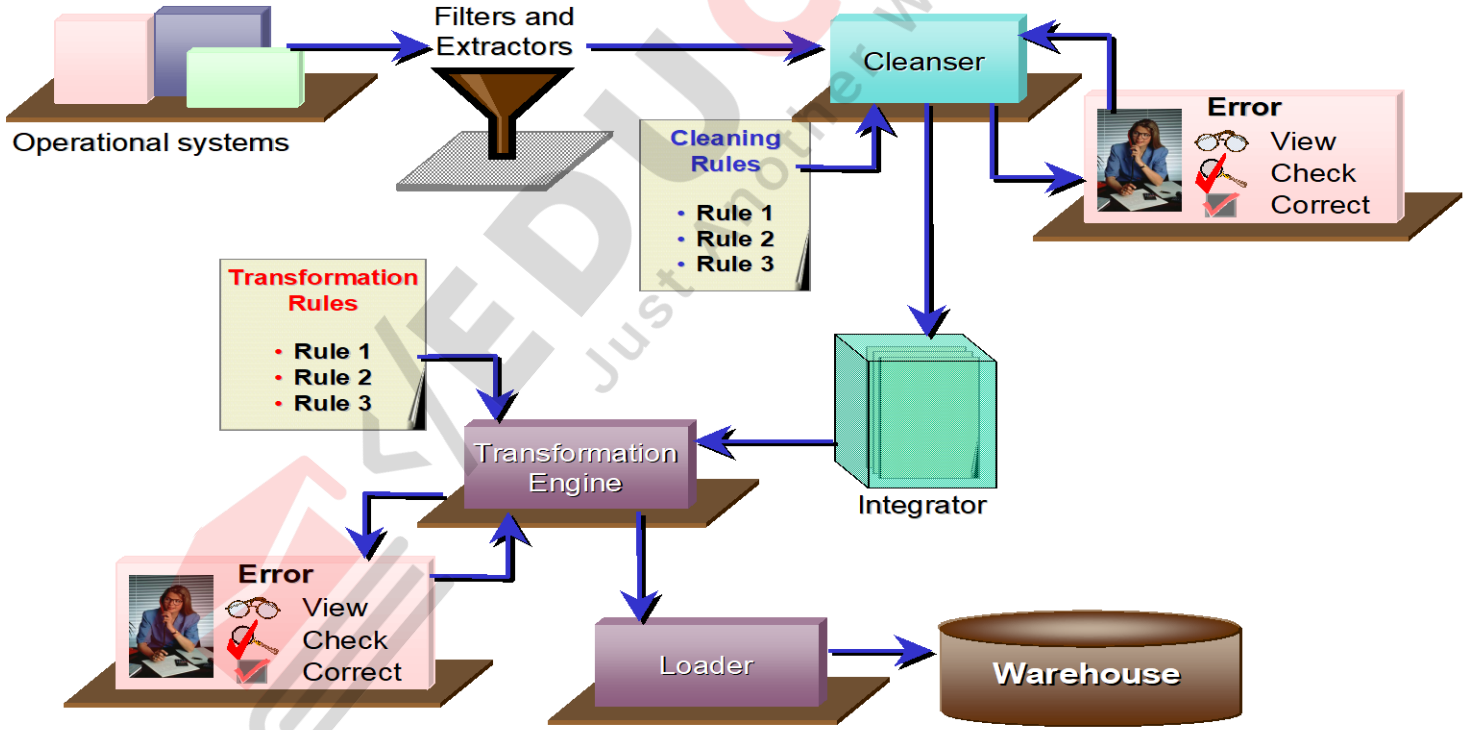

Extraction

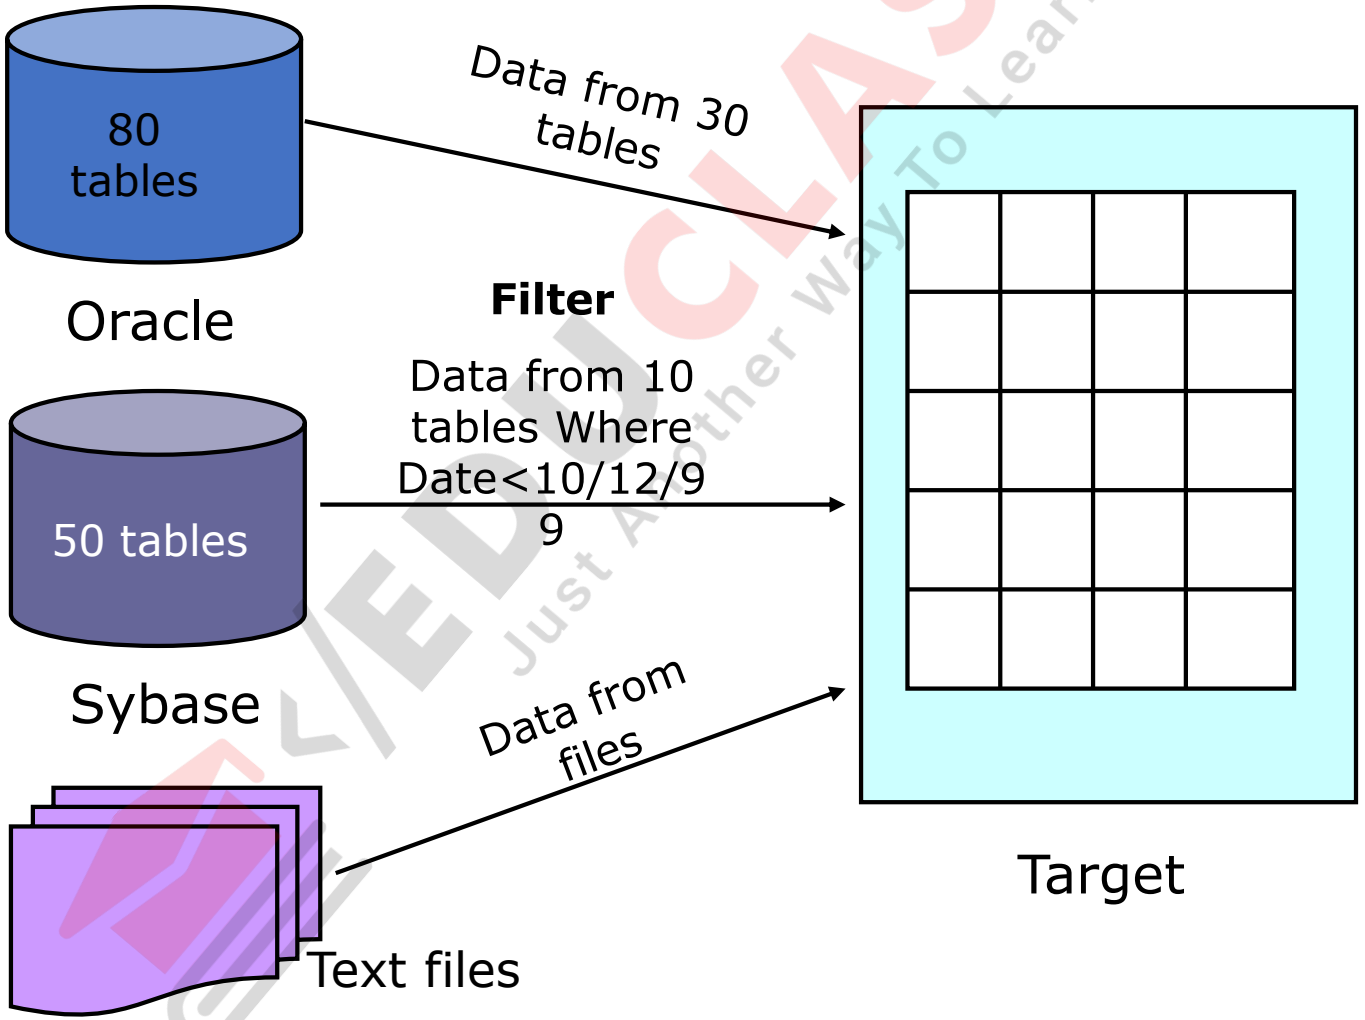

Extraction

- Some of the data elements in any operational database can reasonably be expected to be useful in decision making but others are of less value for that purpose.
- E.g

1. A customer may use an american express card or a personal check to pay for a purchase, but there is probably no reason to store the card number in a datawarehouse.

Extraction

2. In a manufacturing analysis application the customer visible information about products being produced may be relevant, but the internal manufacturing details like which machine ,about component of machine , by which machinist etc are not.

For this reason, it is necessary to extract relevant information from operational database before bringing it into the DW.

Extraction

Many commercial tools are available to help with the extraction process. The user of one of these tools typically has an easy to use windows interface by which to specify the following:

1. Which files or tables are to be accessed in the source database?
2. Which fields are to be extracted from them?(Often done internally via the SQL select

Extraction

statement, though many extraction tools hide the complexity of the SQL from the person defining the extraction.

3. What are those fields to be called in the resulting database?

4. What is the target m/c and database format of the output?

Extraction

5. On what schedule should the extraction be repeated? While initial database loading is a one time operation, regular updates are a common practice.

The input list and the output list of database formats are usually different, since some common formats in transactional databases would not be chosen for use in a dw and vice versa.

Transformation

Types of transformation:

- Structure
- Format
- Conversions
- Classification
- Data Consistency
- Reconciliation of duplicated data

Structural Transformations

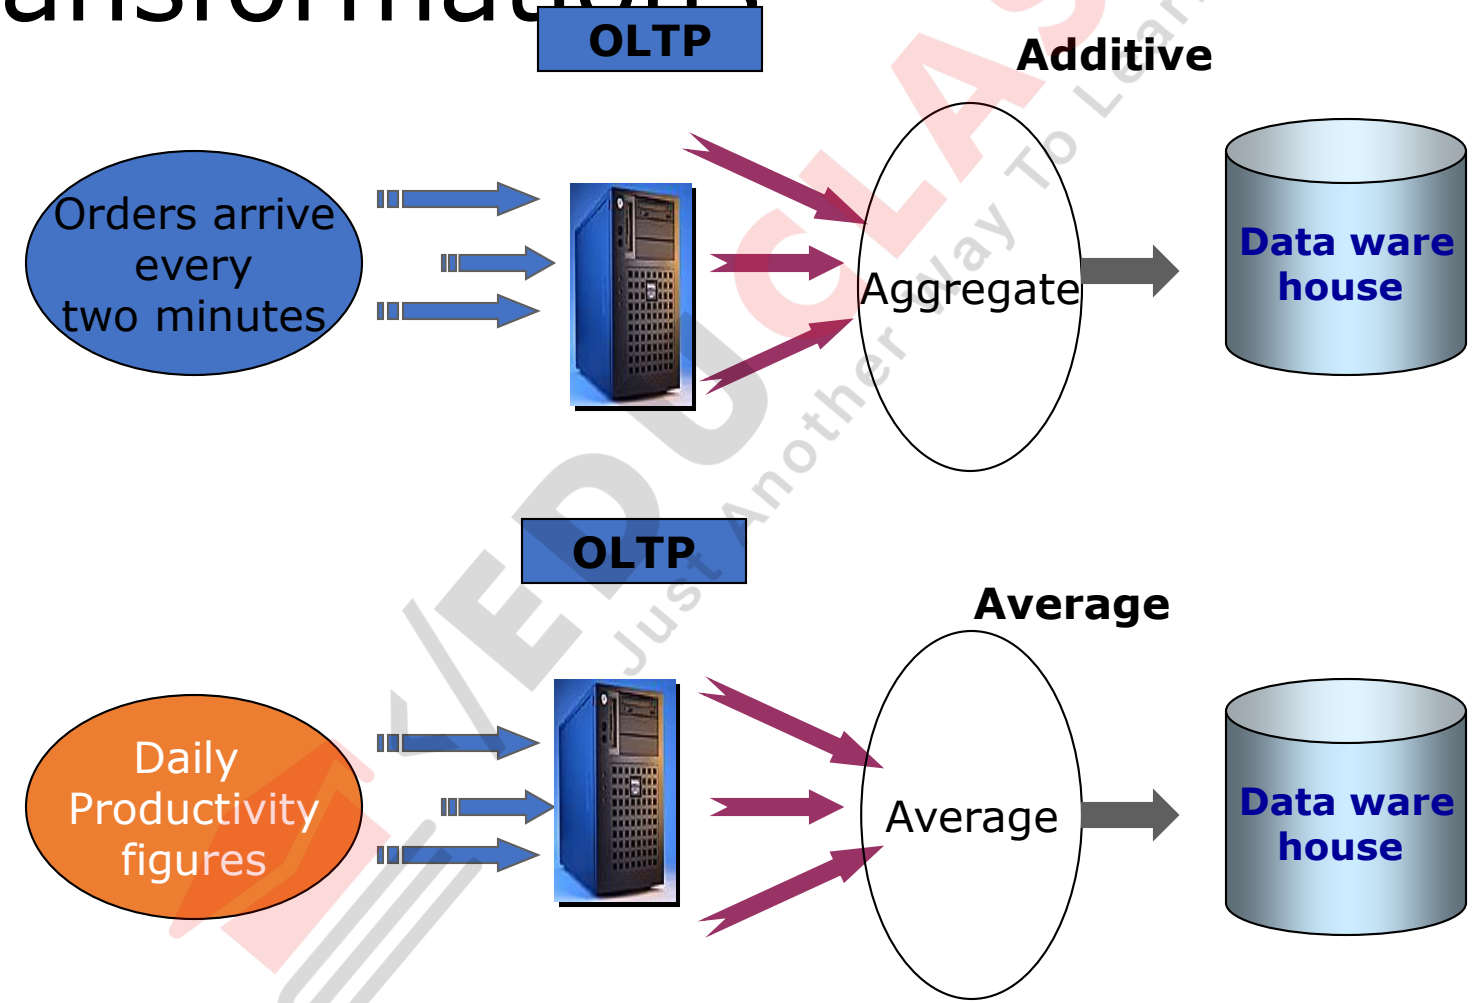

Format Transformation

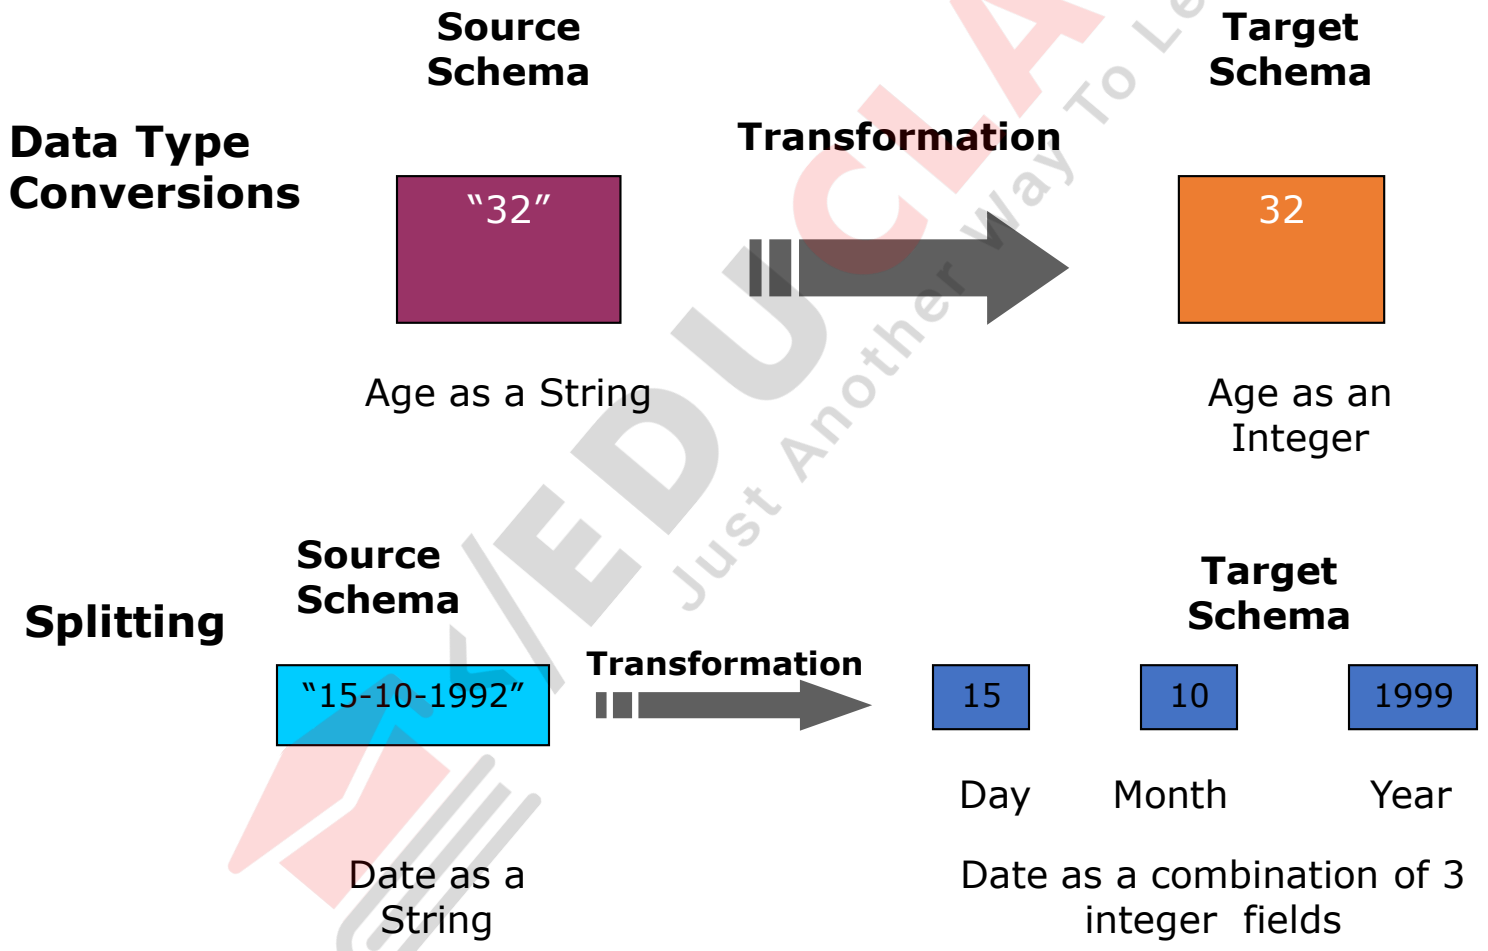

Simple Conversions

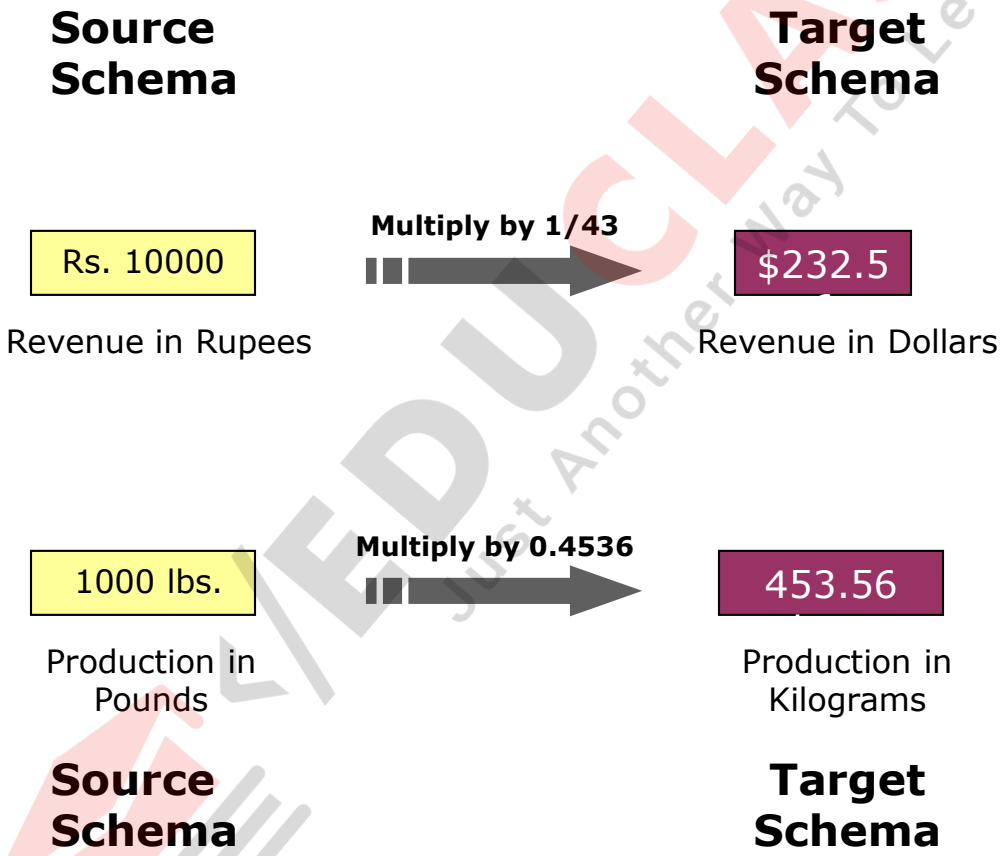

- Transformations using Simple Conversions

Classification

Name	Age
John Black	27
Richard Wayne	53
Jennifer Goldman	45
Helmut Koch	37
Anna Ludwig	32
Shito Maketha	28
Tracy Withman	39
Ada Zhesky	25
David Rosenberg	33
Pankaj Sharma	29
Zhu Ling	44
George Kurtz	27
Rita Hartman	34

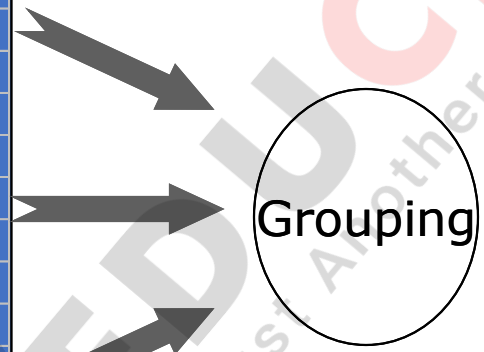

Age Group	Frequency
20-25	1
26-30	4
31-35	3
36-40	2
41-45	2
46-50	1
51-55	1
56-60	0

Data Consistency Transformations

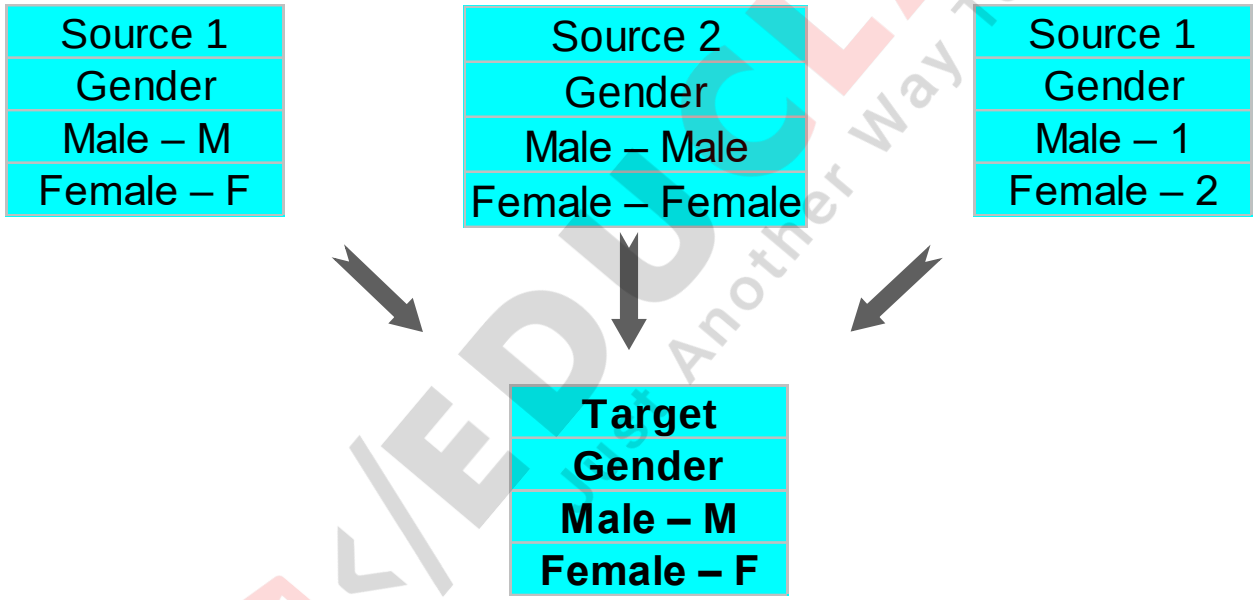

Reconciliation of Duplicated Data

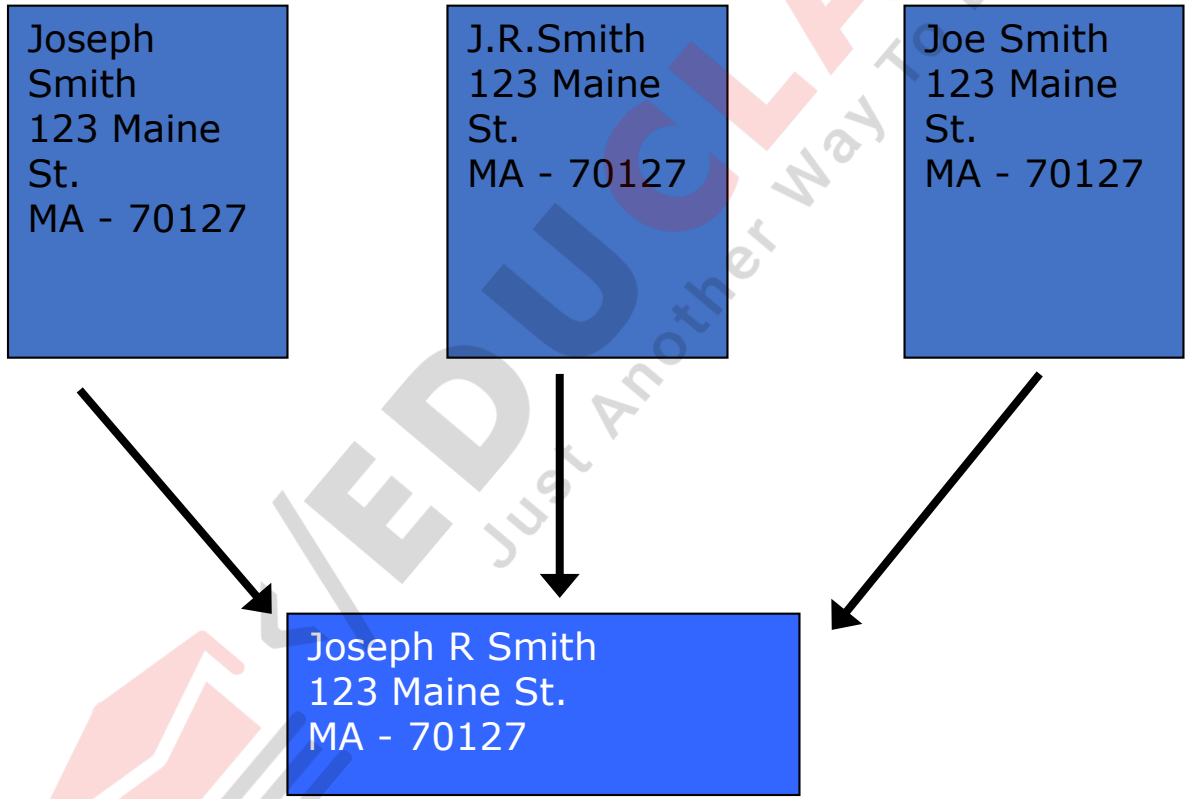

Loading

- After the successful completion of previous

steps , it is finally possible to load the new data into the data warehouse database.

- Loading often implies physical movement of the data from the computer(s) storing the source database(s) to that which will store the data warehouse database, assuming it is different.
- The most common channel for the data movement process is a high-speed communication link.
- High-speed links using local area technology are inexpensive if the two systems are located near each other.

Loading

- Organizations large enough to have them geographically separated are also usually large enough to have high speed wide area links in place.
- Most systems housing operational databases are set up as servers.
- It is then relatively simple matter for the system housing the data warehouse database, or another system that is controlling the load operation, to act as a client and request the necessary data.

Loading

- At the end of the cleansing step , we have data ready to put into the data warehouse database itself.
- Every DBMS normally used to support a DW has an import function that can carry out this process.
- It may be necessary or desirable to close off access to the DW when the loading is taking place.

Loading

- But as the time taken by an incremental update is small and it can take place during periods of little or no DW usage.

DIMENSIONAL MODELING

- A Dimensional Model is a database structure that is optimized for online queries and Data Warehousing tools.
- It is comprised of "fact" and "dimension" tables.

Cont…

- Facts are the measurements/metrics or facts from your business process. For a Sales business process, a measurement would be quarterly sales number
- Dimension provides the context surrounding a business process event. In simple terms, they give who, what, where of a fact. In the Sales business process, for the fact quarterly sales number, dimensions would be
- Who – Customer Names
- Where – Location
- What – Product Name

Dimension and Fact table

1. **Dimension Table**
   Dimension table is one that describe the business entities of an enterprise, represented as hierarchical, categorical information such as time, departments, locations, and products. Dimension tables are sometimes called lookup or reference tables.
2. **Fact Table**
   The centralized table in a star schema is called as FACT table. A fact table typically has two types of columns: those that contain facts and those that are foreign keys to dimension tables. The primary key of a fact table is usually a composite key that is made up of all of its foreign keys.

Cont..

- In the example figure "Sales Dollar" is a fact (measure) and it can be added across several dimensions. Fact tables store different types of measures like additive, non additive and semi additive measures.
- **Measure Types**
- Additive - Measures that can be added across all dimensions.
- Non Additive - Measures that cannot be added across all dimensions.
- Semi Additive - Measures that can be added across few dimensions and not with others.
- Let us use examples to illustrate each of the three types of facts. The first example assumes that we are a retailer, and we have a fact table with the following columns:

Cont…

Date	Store	Product	Sales_Amount

- The purpose of this table is to record the sales amount for each product in each store on a daily basis.
- **Sales_Amount** is the fact.
- In this case, **Sales_Amount** is an additive fact, because you can sum up this fact along any of the three dimensions present in the fact table -- date, store, and product.
- For example, the sum of **Sales_Amount** for all 7 days in a week represent the total sales amount for that week.
- A bank with the following fact table:

Date	Account	Current_Balance	Profit_Margin

- The purpose of this table is to record the current balance for each account at the end of each day, as well as the profit margin for each account for each day.

Cont..

- **Current_Balance** and **Profit_Margin** are the facts.
- **Current_Balance** is a semi-additive fact, as it makes sense to add them up for all accounts (what's the total current balance for all accounts in the bank?).
- But it does not make sense to add them up through time (adding up all current balances for a given account for each day of the month does not give us any useful information).
- **Profit_Margin** is a non-additive fact, for it does not make sense to add them up for the account level or the day level.

Cont..

- In the real world, it is possible to have a fact table that contains no measures or facts. These tables are called as **Factless Fact** tables.
- Eg: A fact table which has only product key and date key is a factless fact. There are no measures in this table. But still you can get the number products sold over a period of time.

Star Schema

- **Star Schema** in data warehouse, in which the center of the star can have one fact table and a number of associated dimension tables.
- It is known as star schema as its structure resembles a star.
- The Star Schema data model is the simplest type of Data Warehouse schema.
- It is also known as Star Join Schema and is optimized for querying large data sets.
- In the following Star Schema example, the fact table is at the center which contains keys to every dimension table like Dealer_ID, Model ID, Date_ID, Product_ID, Branch_ID & other attributes like Units sold and revenue.

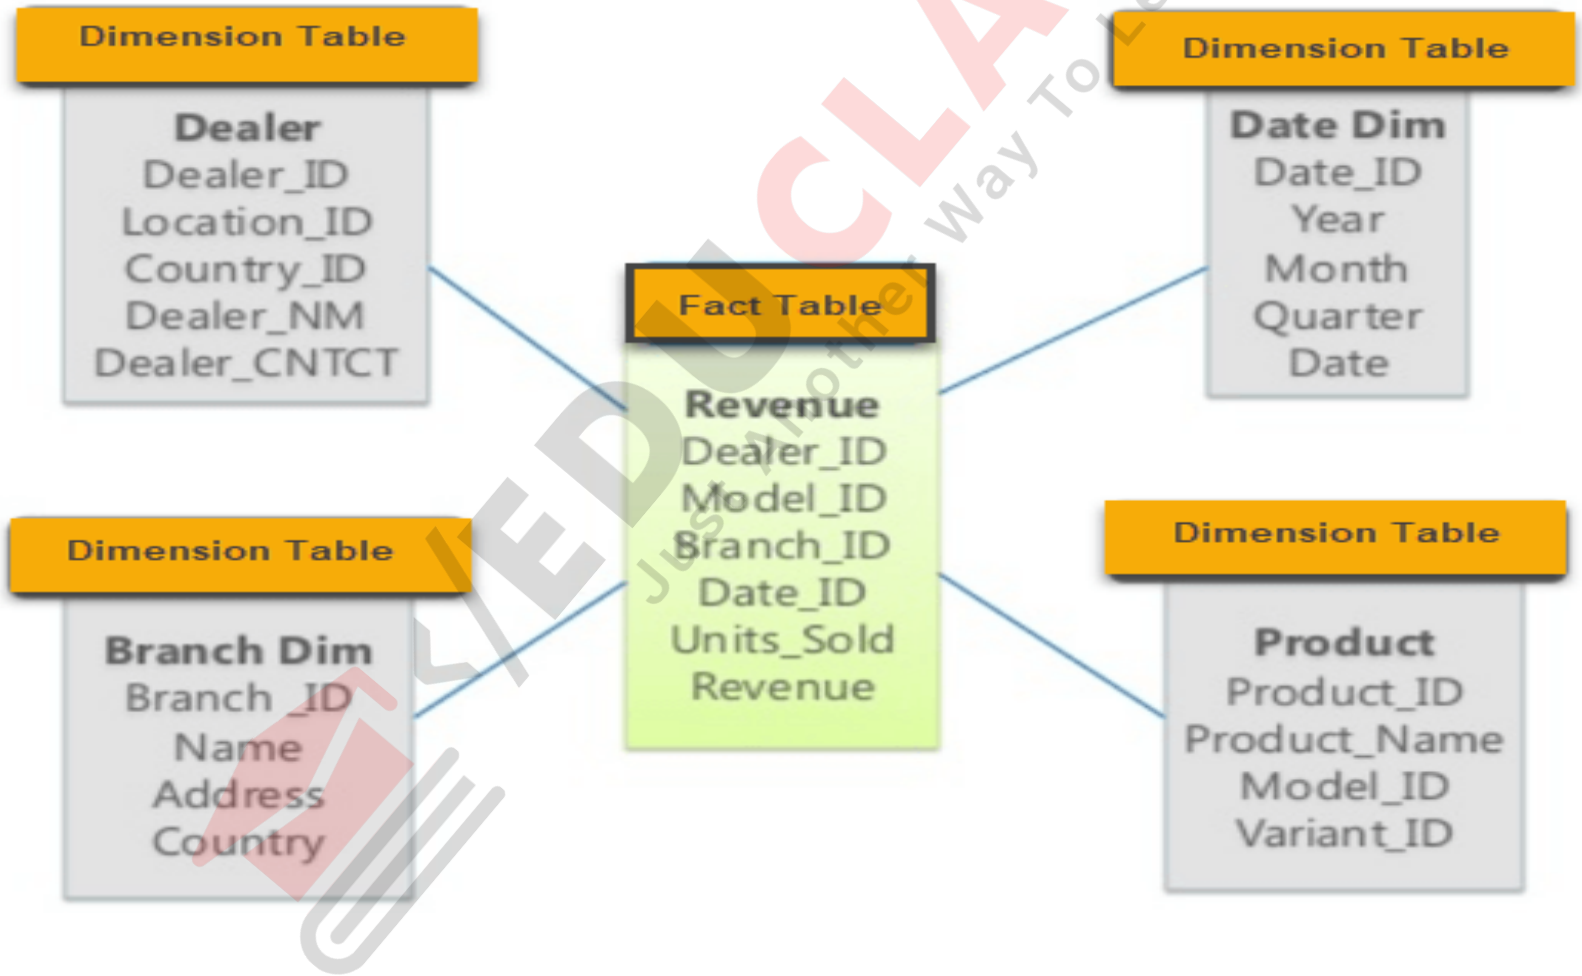

Cont…

- **Characteristics of Star Schema:**
- Every dimension in a star schema is represented with the only one-dimension table.
- The dimension table should contain the set of attributes.
- The dimension table is joined to the fact table using a foreign key
- The dimension table are not joined to each other
- Fact table would contain key and measure
- The Star schema is easy to understand and provides optimal disk usage.
- The dimension tables are not normalized. For instance, in the above figure, Country_ID does not have Country lookup table as an OLTP design would have.
- The schema is widely supported by BI Tools

Snowflake Schema

- **Snowflake Schema** in data warehouse is a logical arrangement of tables in a multidimensional database such that the ER diagram resembles a snowflake shape. A Snowflake Schema is an extension of a Star Schema, and it adds additional dimensions. The dimension tables are normalized which splits data into additional tables.
- In the following Snowflake Schema example, Country is further normalized into an individual table.

Cont…

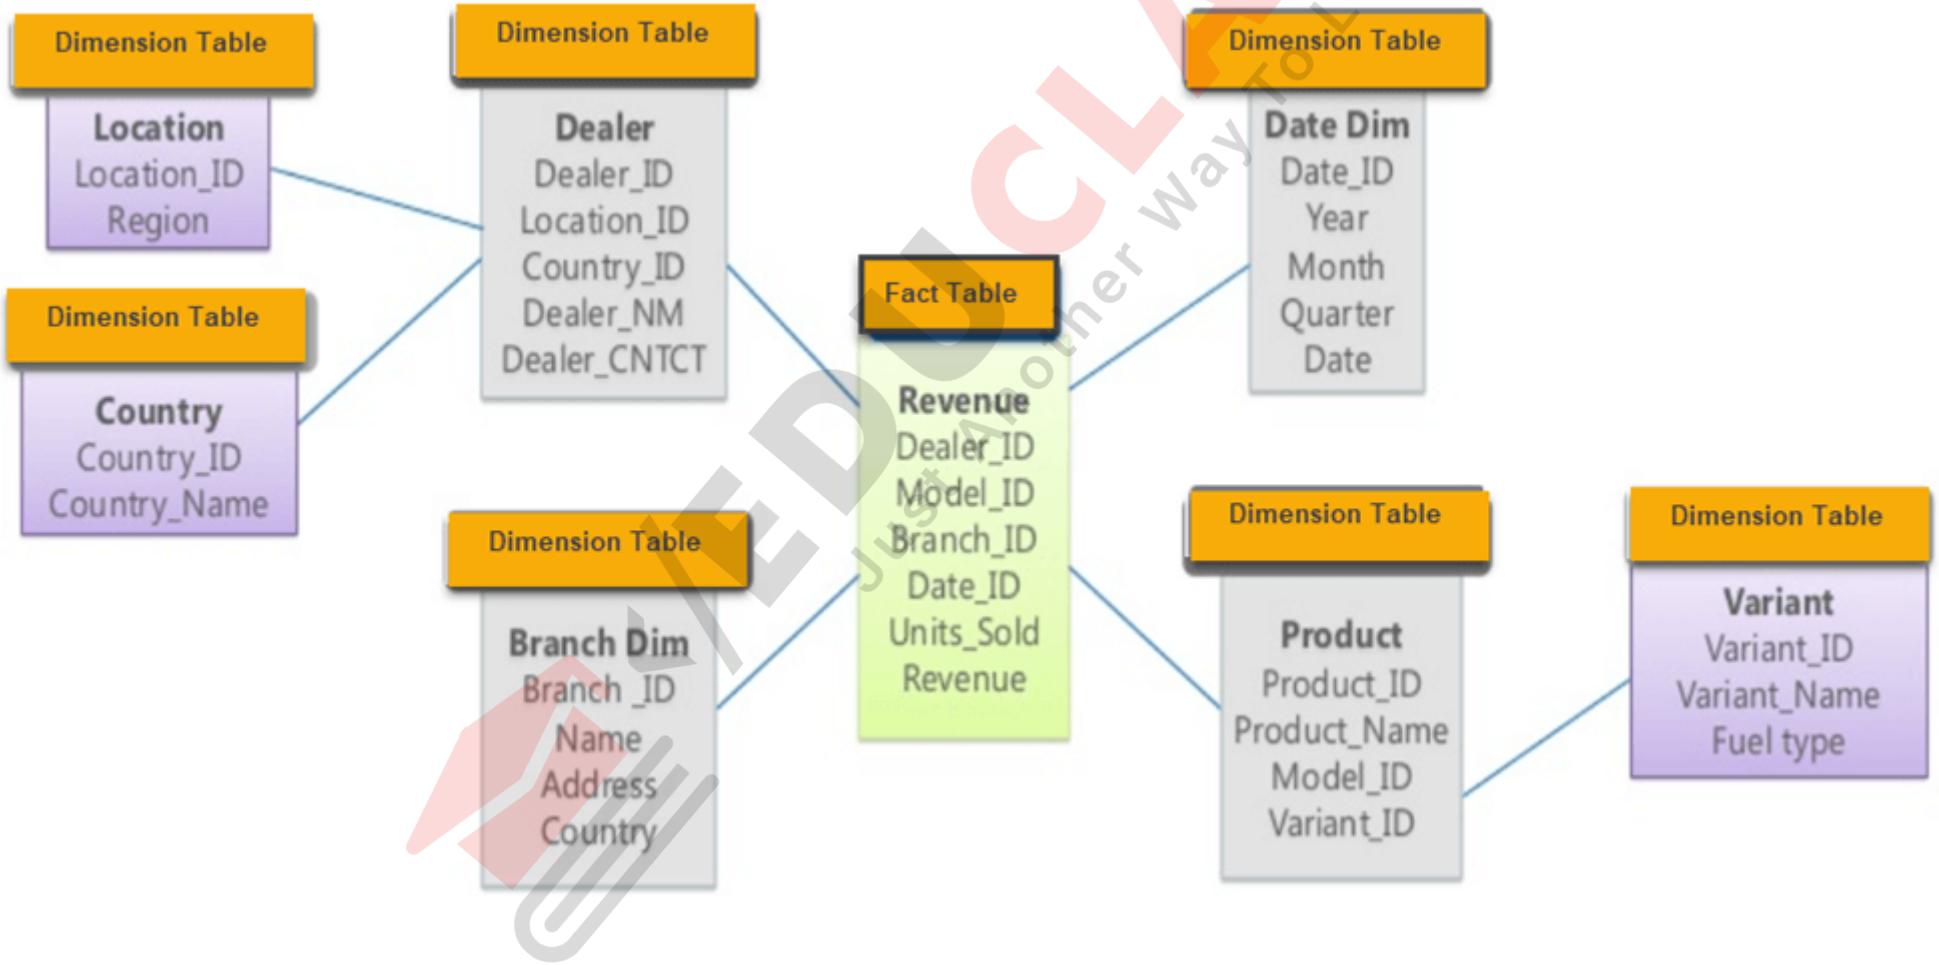

Cont…

- **Characteristics of Snowflake Schema:**
- The main benefit of the snowflake schema it uses smaller disk space.
- Easier to implement a dimension is added to the Schema
- Due to multiple tables query performance is reduced
- The primary challenge that you will face while using the snowflake Schema is that you need to perform more maintenance efforts because of the more lookup tables.

Star Schema	Snowflake Schema
Hierarchies for the dimensions are stored in the dimensional table.	Hierarchies are divided into separate tables.
It contains a fact table surrounded by dimension tables.	One fact table surrounded by dimension table which are in turn surrounded by dimension table
In a star schema, only single join creates the relationship between the fact table and any dimension tables.	A snowflake schema requires many joins to fetch the data.
Simple DB Design.	Very Complex DB Design.
Denormalized Data structure and query also run faster.	Normalized Data Structure.
High level of Data redundancy	Very low-level data redundancy
Single Dimension table contains aggregated data.	Data Split into different Dimension Tables.
Cube processing is faster.	Cube processing might be slow because of the complex join.

Fact Constellation Schema

- A **Galaxy Schema** contains two fact table that share dimension tables between them. It is also called Fact Constellation Schema.
- The schema is viewed as a collection of stars hence the name Galaxy Schema.

Cont..

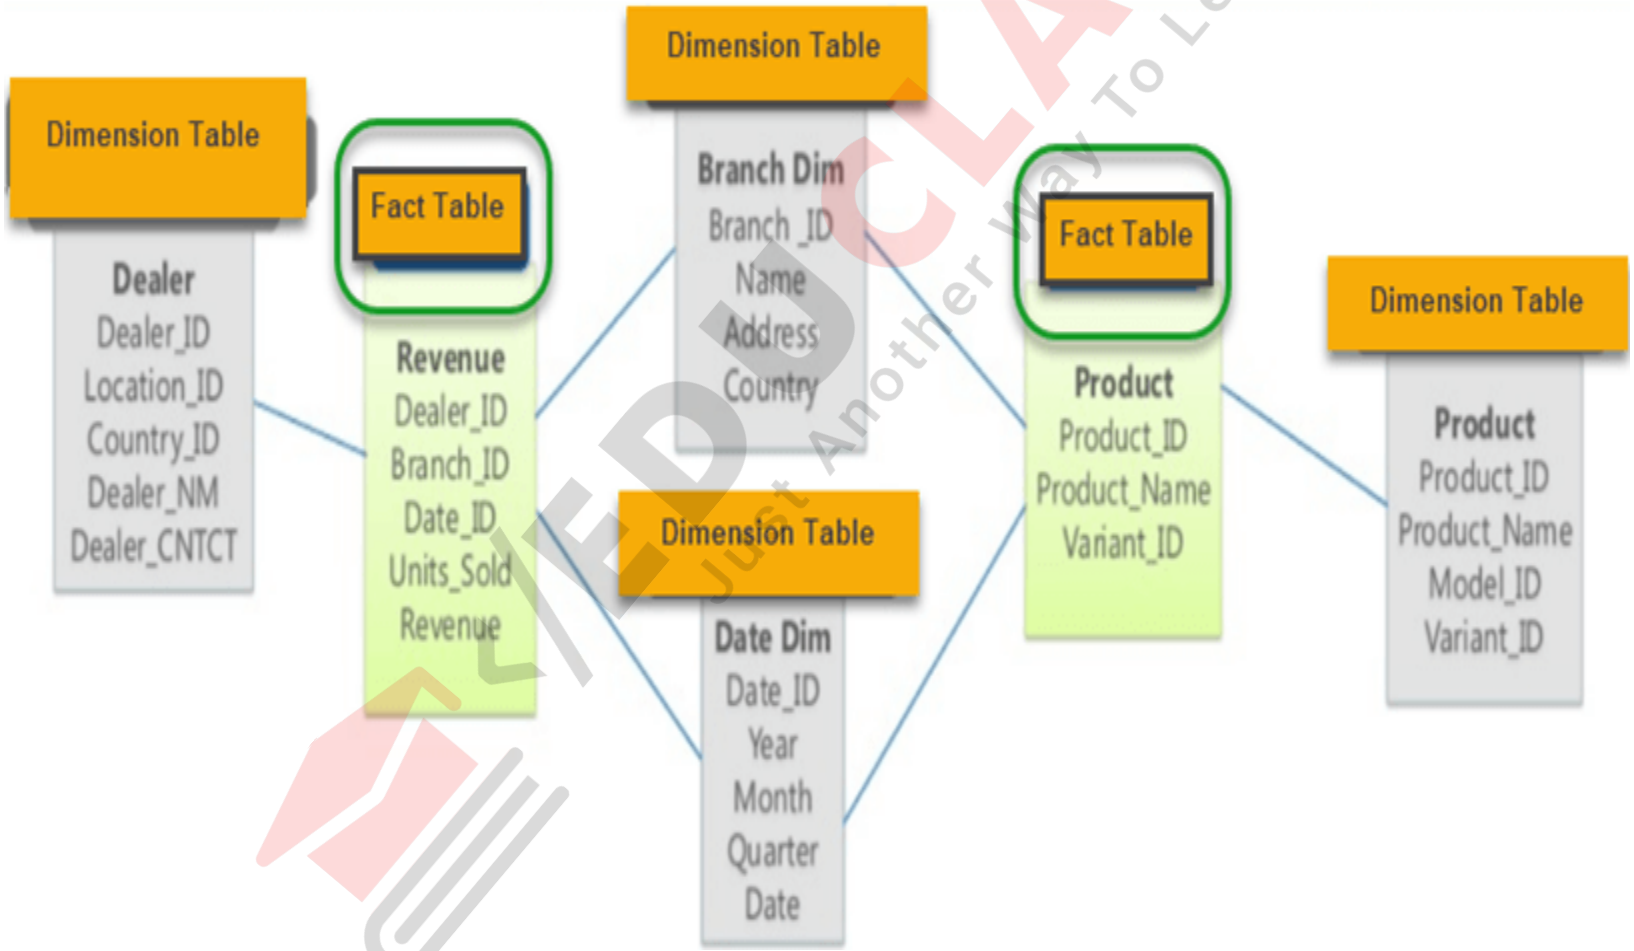

Cont...

- The dimensions in this schema are separated into separate dimensions based on the various levels of hierarchy.
- For example, if geography has four levels of hierarchy like region, country, state, and city then Galaxy schema should have four dimensions.
- Moreover, it is possible to build this type of schema by splitting the one-star schema into more Star schemes.
- The dimensions are large in this schema which is needed to build based on the levels of hierarchy.
- This schema is helpful for aggregating fact tables for better understanding.

OLAP Architecture

- Designed to use both operational and data warehouse data
- Defined as an “advanced data analysis environment that supports decision making, business modeling, and an operation’s research activities”
- In most implementations, data warehouse and OLAP are interrelated and complementary environments

FIGURE 13.9 OLAP server with local mini data marts

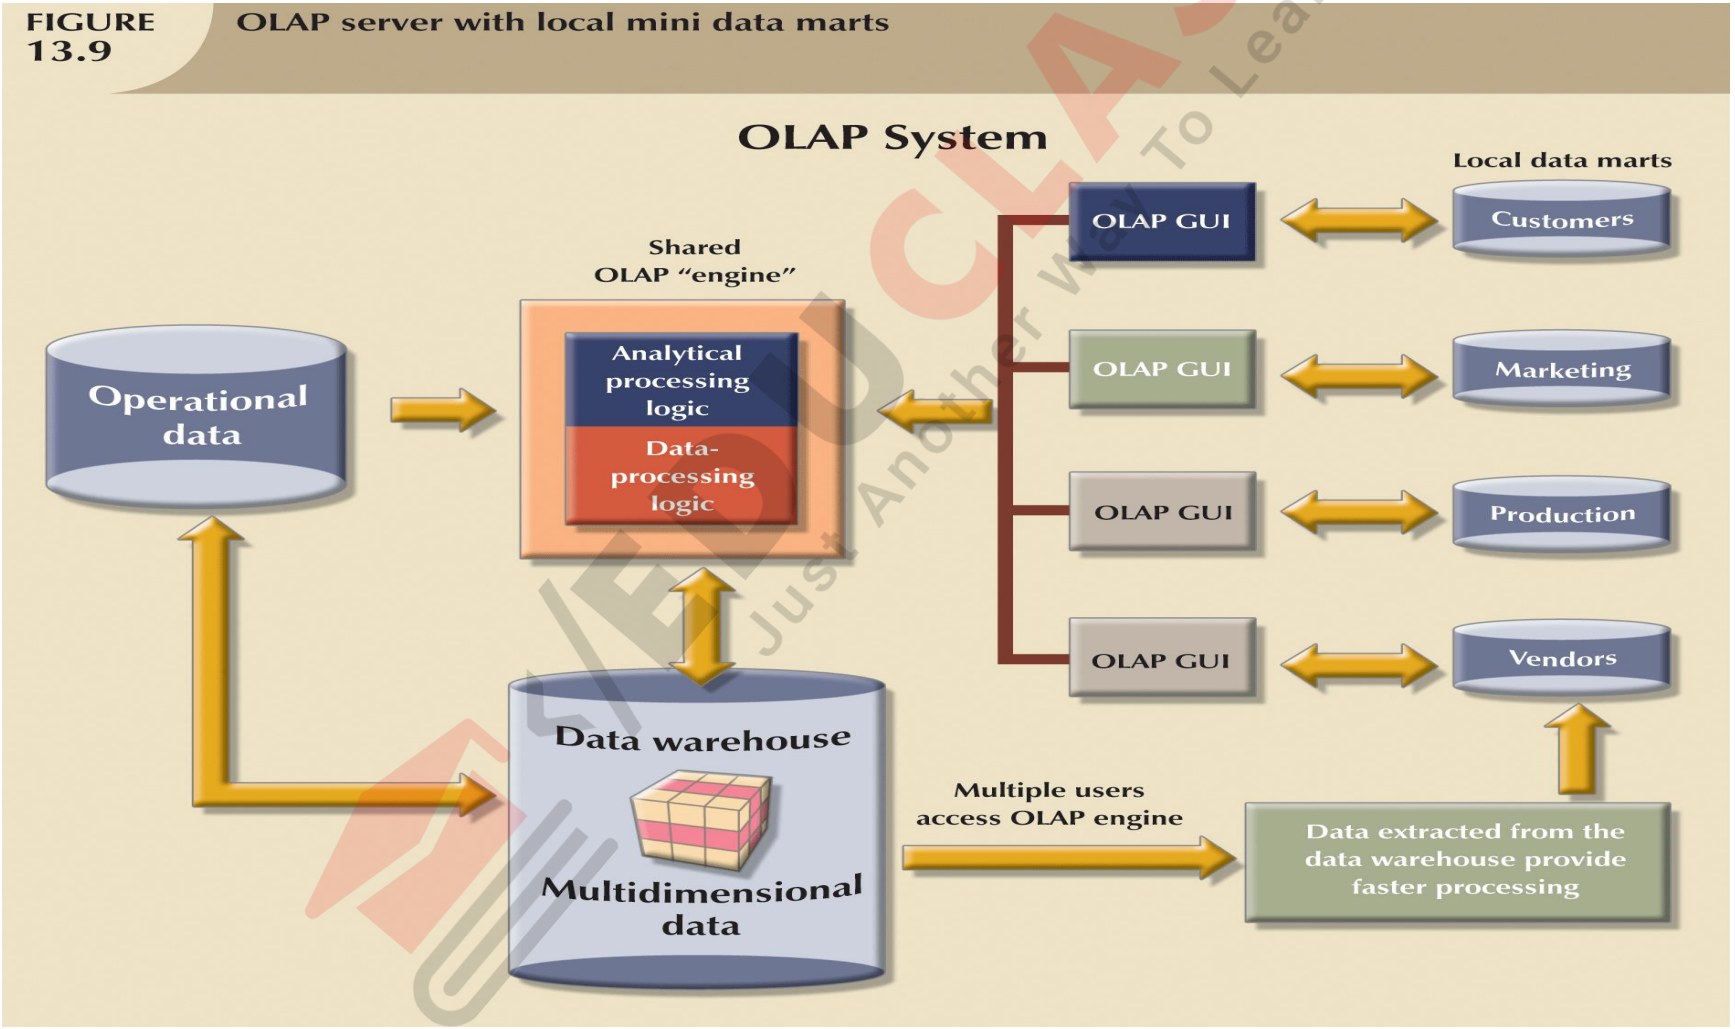

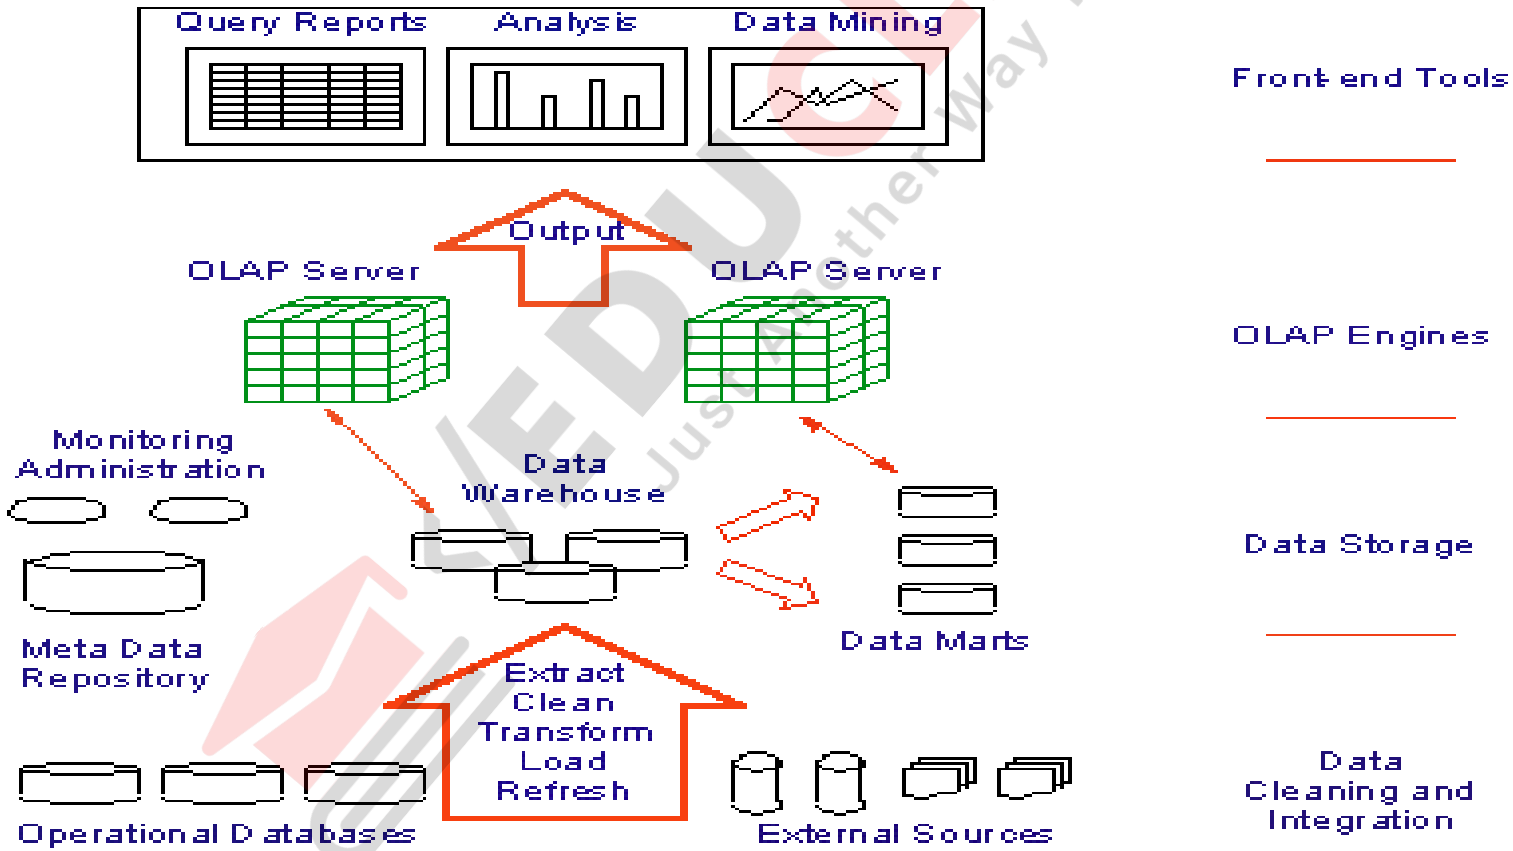

OLAP : Multidimensional data model

- MOLAP : Multidimensional OLAP

➢ Utilizes a pre-calculated dataset, commonly referred to as a data cube, that contains all the possible answers to a given range of questions.

➢ Very fast response.

➢ Ability to quickly write back data into the data set.

➢ Limited scalability

➢ Inability to contain detailed data.

OLAP : Multidimensional data model

- ROLAP : Relational OLAP

➢ OLAP systems that store all information (dimension and fact tables) as relations are called relational OLAP.

➢ Do not use pre-calculated data cubes.

➢ Intercept the query and pose the question to the standard relational database and its tables in order to bring back the data required to answer the question.

MOLAP vs ROLAP

Sr.No.	MOLAP	ROLAP
1	Information retrieval is fast.	Information retrieval is comparatively slow.
2	Uses sparse array to store data-sets.	Uses relational table.
3	MOLAP is best suited for inexperienced users, since it is very easy to use.	ROLAP is best suited for experienced users.
4	Maintains a separate database for data cubes.	It may not require space other than available in the Data warehouse.
5	DBMS facility is weak.	DBMS facility is strong.

ON-LINE ANALYTICAL PROCESSING (OLAP)

OLAP can be defined in five words – Fast Analysis of Shared Multi-dimensional Information(FASMI) .

- **Fast** : means that the system is targeted to deliver most responses to users within about 5 seconds, with the simplest analysis not taking more than one second and very few taking more than 20 seconds.
- **Analysis**: means that the system can cope with any business logic and statistical analysis that is relevant for the application and the user, and keep it easy enough for the target user.
- **Shared**: means that the system implements all the security requirements for confidentiality and if multiple write access is needed, concurrent update locking at an appropriate level.
- **Multi-dimensional**: means that the system must provide a multi-dimensional conceptual view of the data, including full support for hierarchies and multiple hierarchies.
- **Information**: is refined data that is accurate, timely and relevant to the user.

OLAP Operations

- OLAP products must provide at least the following minimal set of functions:

➢1. ROLL-UP: The roll-up operation collapses the dimension hierarchy along a particular dimension(s) so as to present the remaining dimensions at a coarser level of granularity.

  - E.g., Given total sales by city, we can roll-up to get sales by state.

OLAP Operations

- 2. DRILL-DOWN: In contrast, the drill-down function allows users to obtain a more detailed view of a given dimension.
  - E.g., Given total sales by state, can drill-down to get total sales by city.
  - E.g., Can also drill-down on different dimension to get total sales by product for each state.

3. SLICE: The objective is to extract a slice of the original cube corresponding to a single value of a given dimension. No aggregation is required with this option.

Slicing-Wireless Mouse

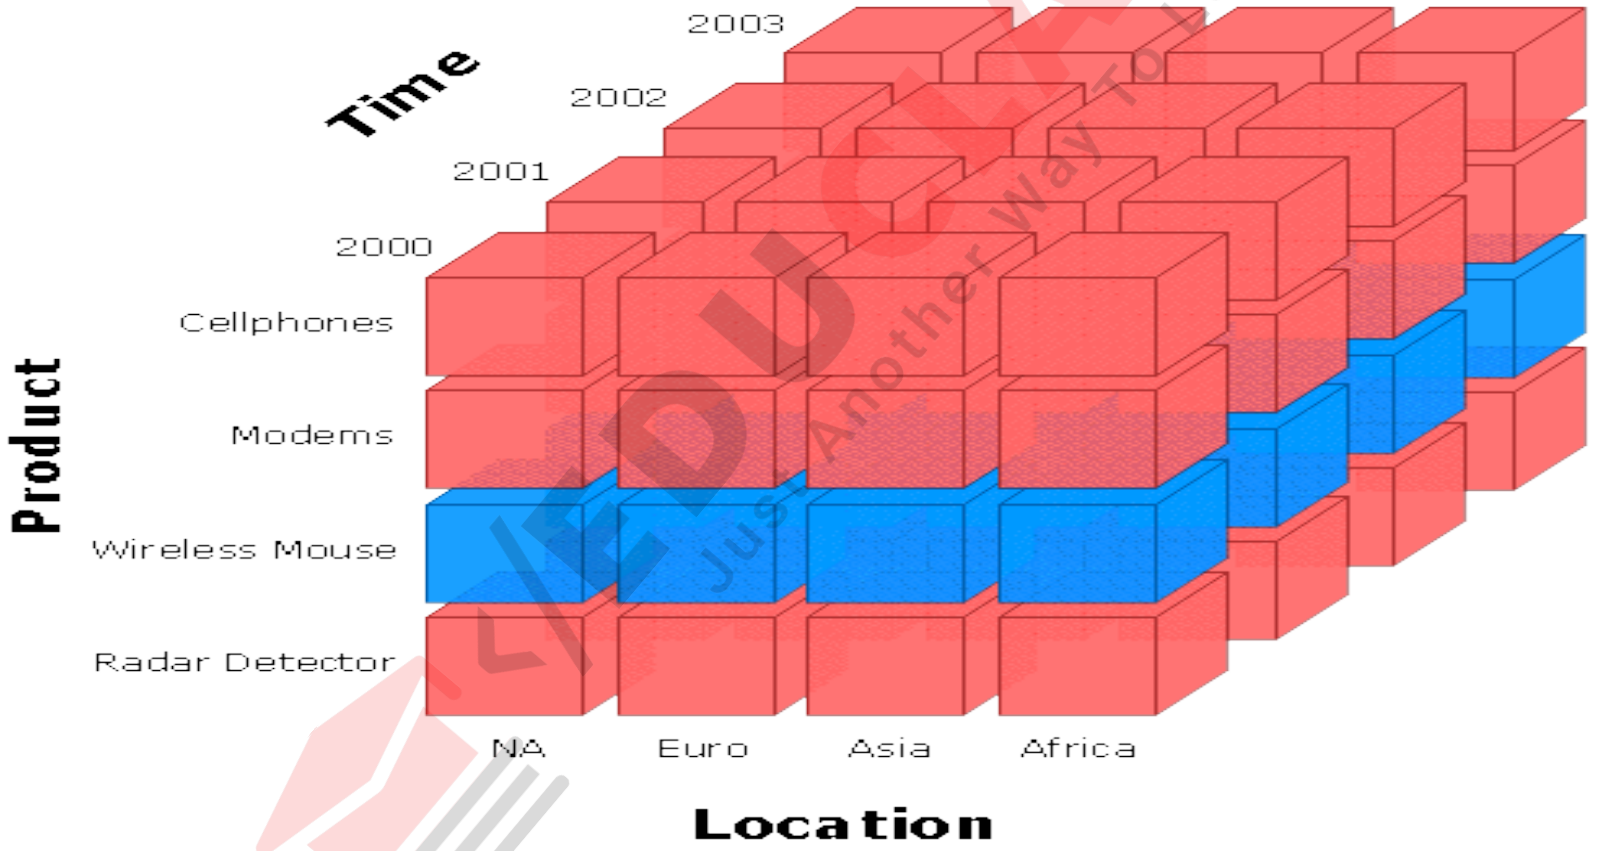

Slicing-Asia

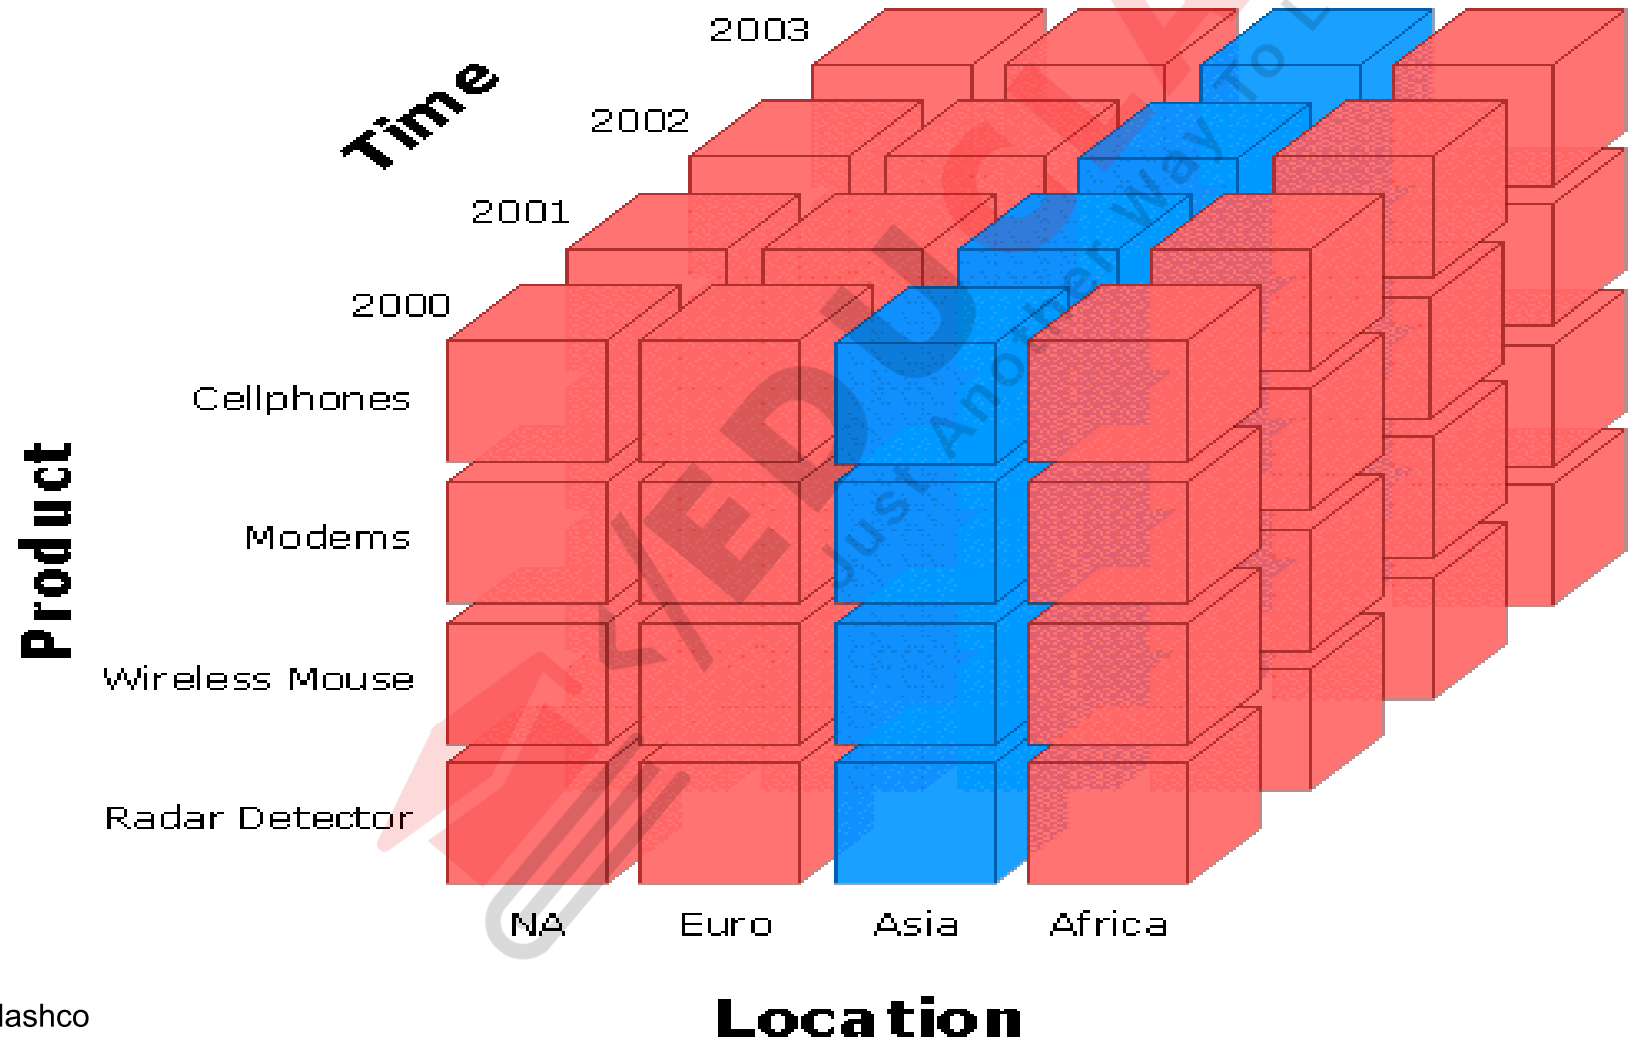

OLAP Operations

- 4. DICE: A related operation is the dice. In this case, we are defining a subcube of the original space. In other words, by specifying value ranges on one or more dimensions, the user can highlight meaningful blocks of aggregated data.

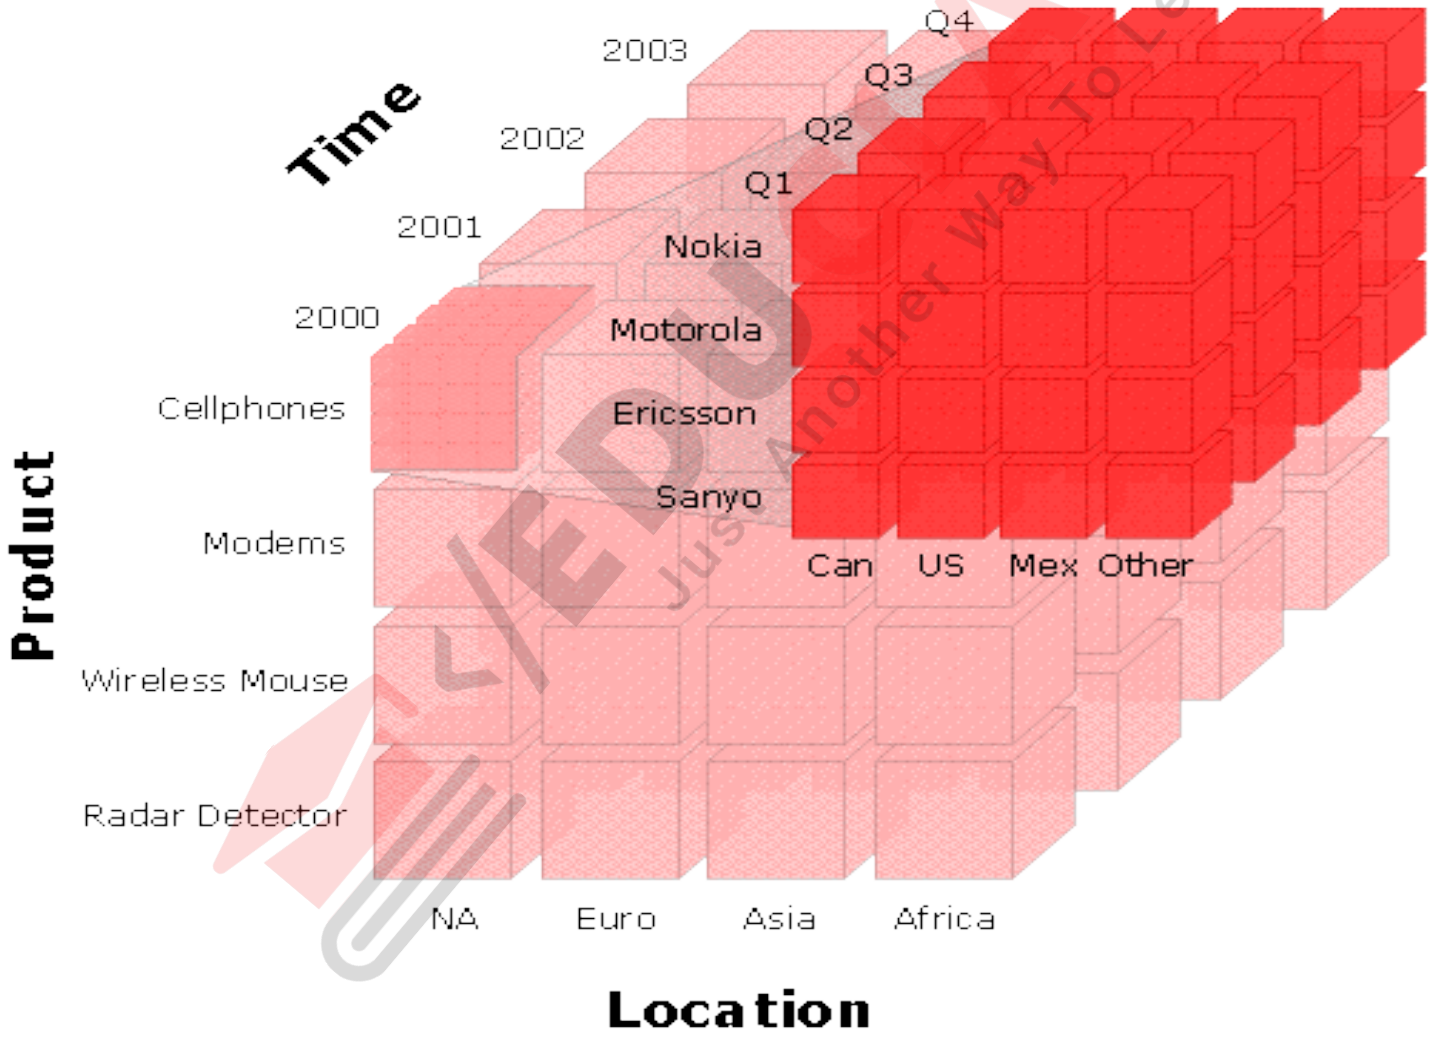

OLAP Operations

- 5. PIVOT: The pivot is a simple but effective operation that allows OLAP users to visualize cube values in more natural and intuitive ways.
- Consider a tabular presentation of the Sales table. If we pivot on the Location and Time dimensions, we obtain a table of total sales for each location for each time value.
- The result of pivoting, called a cross-tabulation.

OLAP Operations

- E.g., Pivoting on Location and Time yields this **cross-tabulation**

	WI	CA	Total
1995	63	81	144
1996	38	107	145
1997	75	35	110
Total	176	223	399

Figure 25.5 Cross-Tabulation of Sales by Year and St

Pivoting

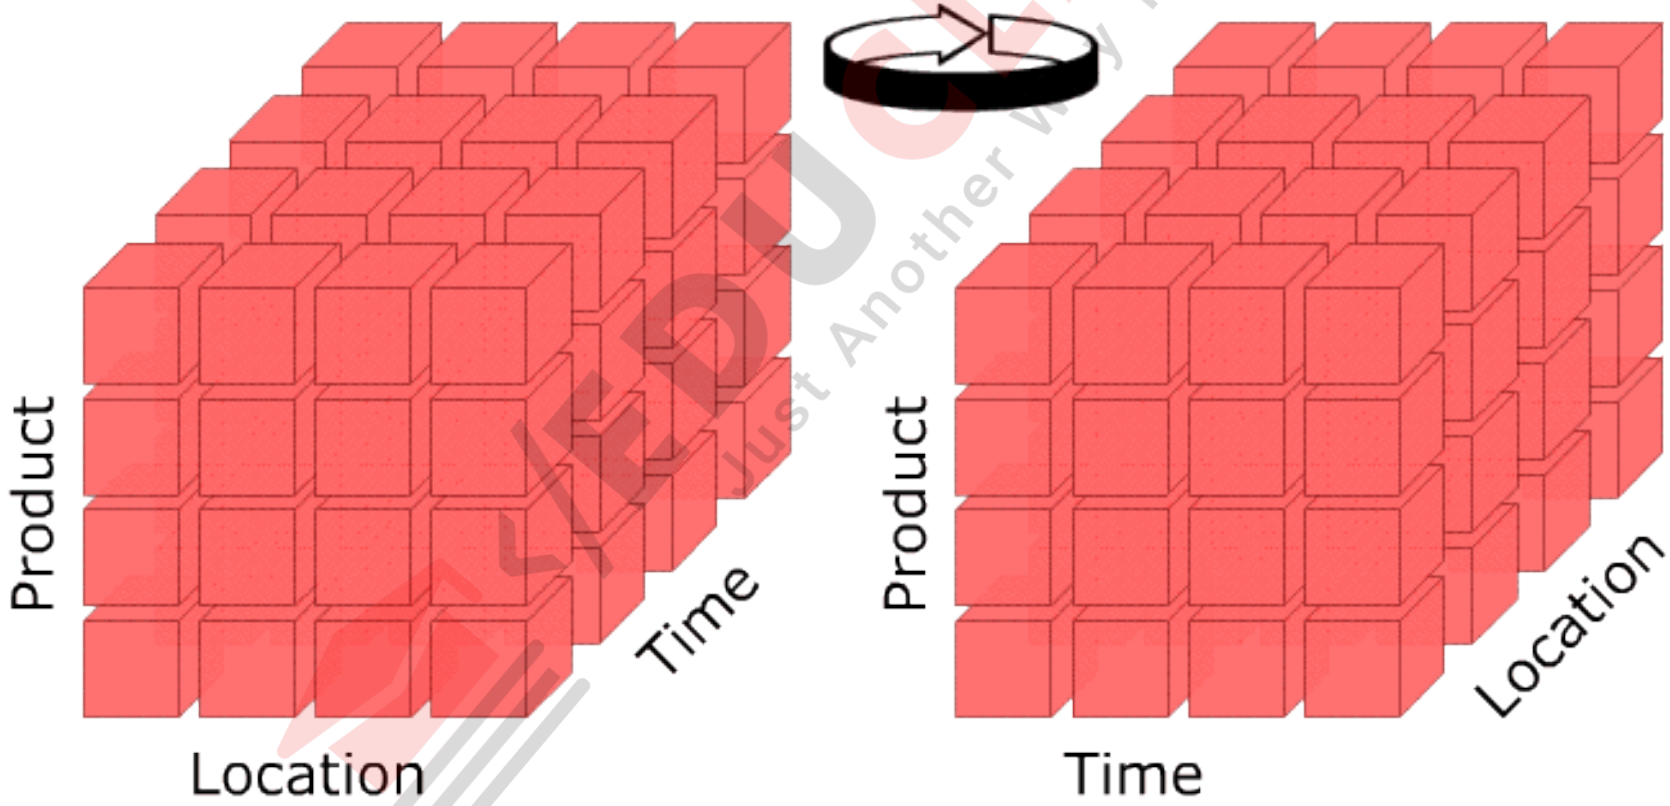

OLAP vs OLTP

Sr.No.	Data Warehouse (OLAP)	Operational Database (OLTP)
1	Involves historical processing of information.	Involves day-to-day processing.
2	OLAP systems are used by knowledge workers such as executives, managers and analysts.	OLTP systems are used by clerks, DBAs, or database professionals.
3	Useful in analyzing the business.	Useful in running the business.
4	It focuses on Information out.	It focuses on Data in.
5	Based on Star Schema, Snowflake, Schema and Fact Constellation Schema.	Based on Entity Relationship Model.
6	Contains historical data.	Contains current data.
7	Provides summarized and consolidated data.	Provides primitive and highly detailed data.
8	Provides summarized and multidimensional view of data.	Provides detailed and flat relational view of data.
9	Number or users is in hundreds.	Number of users is in thousands.
10	Number of records accessed is in millions.	Number of records accessed is in tens.
11	Database size is from 100 GB to 1 TB	Database size is from 100 MB to 1 GB.
12	Highly flexible.	Provides high performance.