

Материализованная интеграция данных и организация хранилищ больших данных

Сергей Ступников
ФИЦ ИУ РАН
sstupnikov@ipiran.ru

Материалы к курсу

- <https://synthesis.frccsc.ru/synthesis/student/BigData/course-data-warehousing.html>
- synthesis.frccsc.ru ->
- For MSU Students ->
- Материализованная интеграция данных и организация хранилищ больших данных

Отчетность

- Посещение 20/100
- Контрольные 20/100
- Лабораторные работы 20/100
- Экзамен 40/100

Оценки за курс:

- 85 – «5»
- 70 – «4»
- 55 – «3»

Лабораторные работы

- Выбор OLAP-сервера (поддержка MDX?)
 - Mondrian (<https://github.com/pentaho/mondrian>)
 - eMondrian (<https://sergeisemenkov.github.io/eMondrian/docs/Introduction>)
 - ???
- Выбор предметной области, схема хранилища
- Выбор источников данных
- Реализация ETL для источников данных
- Развертывание OLAP-куба, реализация запросов
- ❑ Дополнительные баллы за подробное описание

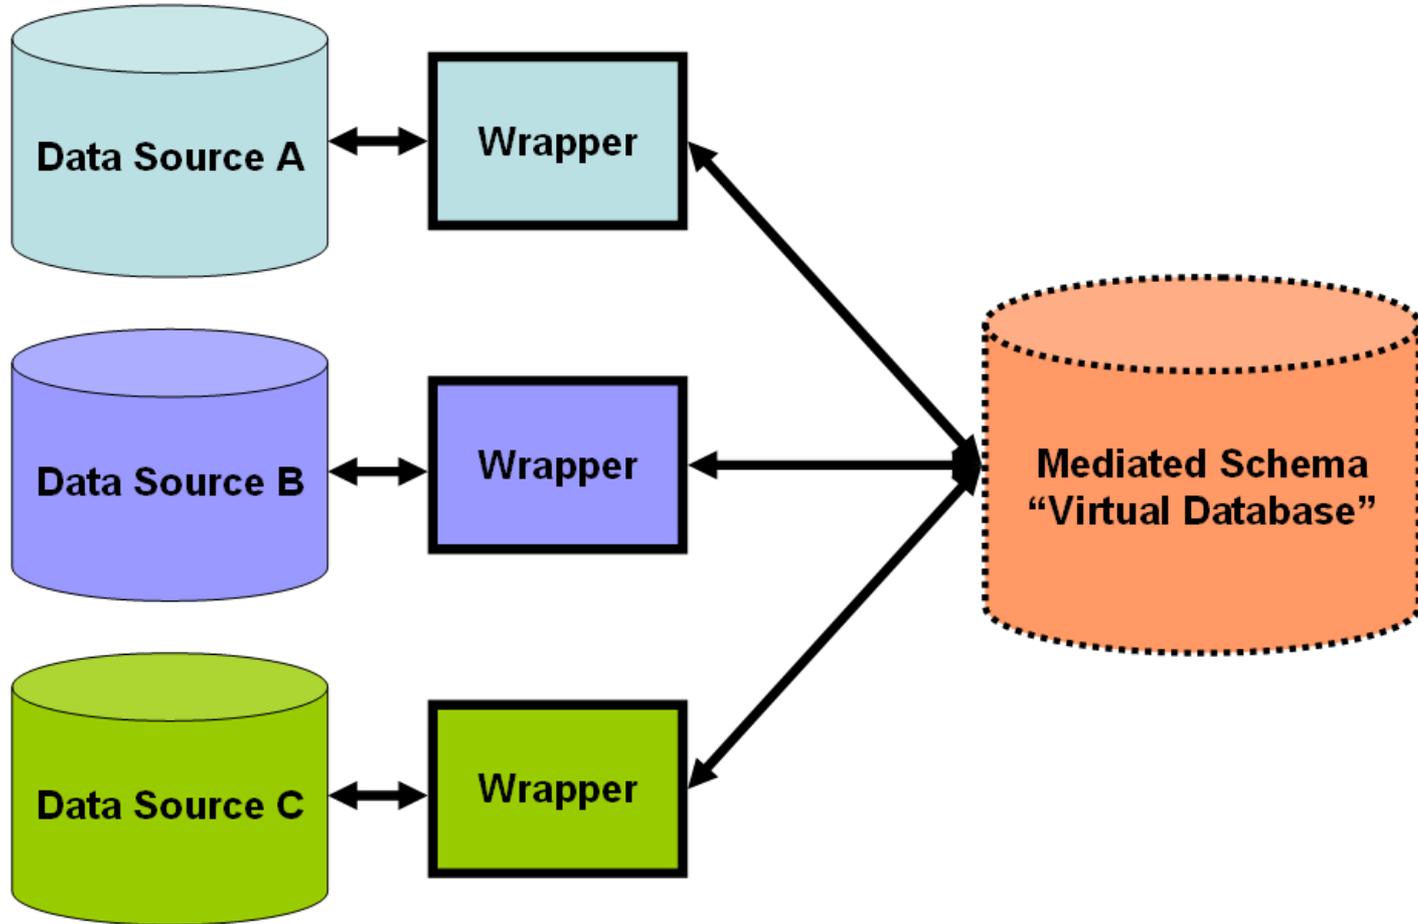
Course Outline

- Теоретические основания материализованной интеграции. Data Exchange
- Архитектуры хранилищ данных
- Проектирование хранилищ данных
 - Anchor Modeling
- Аналитика над хранилищами данных, OLAP
- Поддержка ETL-преобразований
- Построение хранилищ данных на основе параллельных машин баз данных

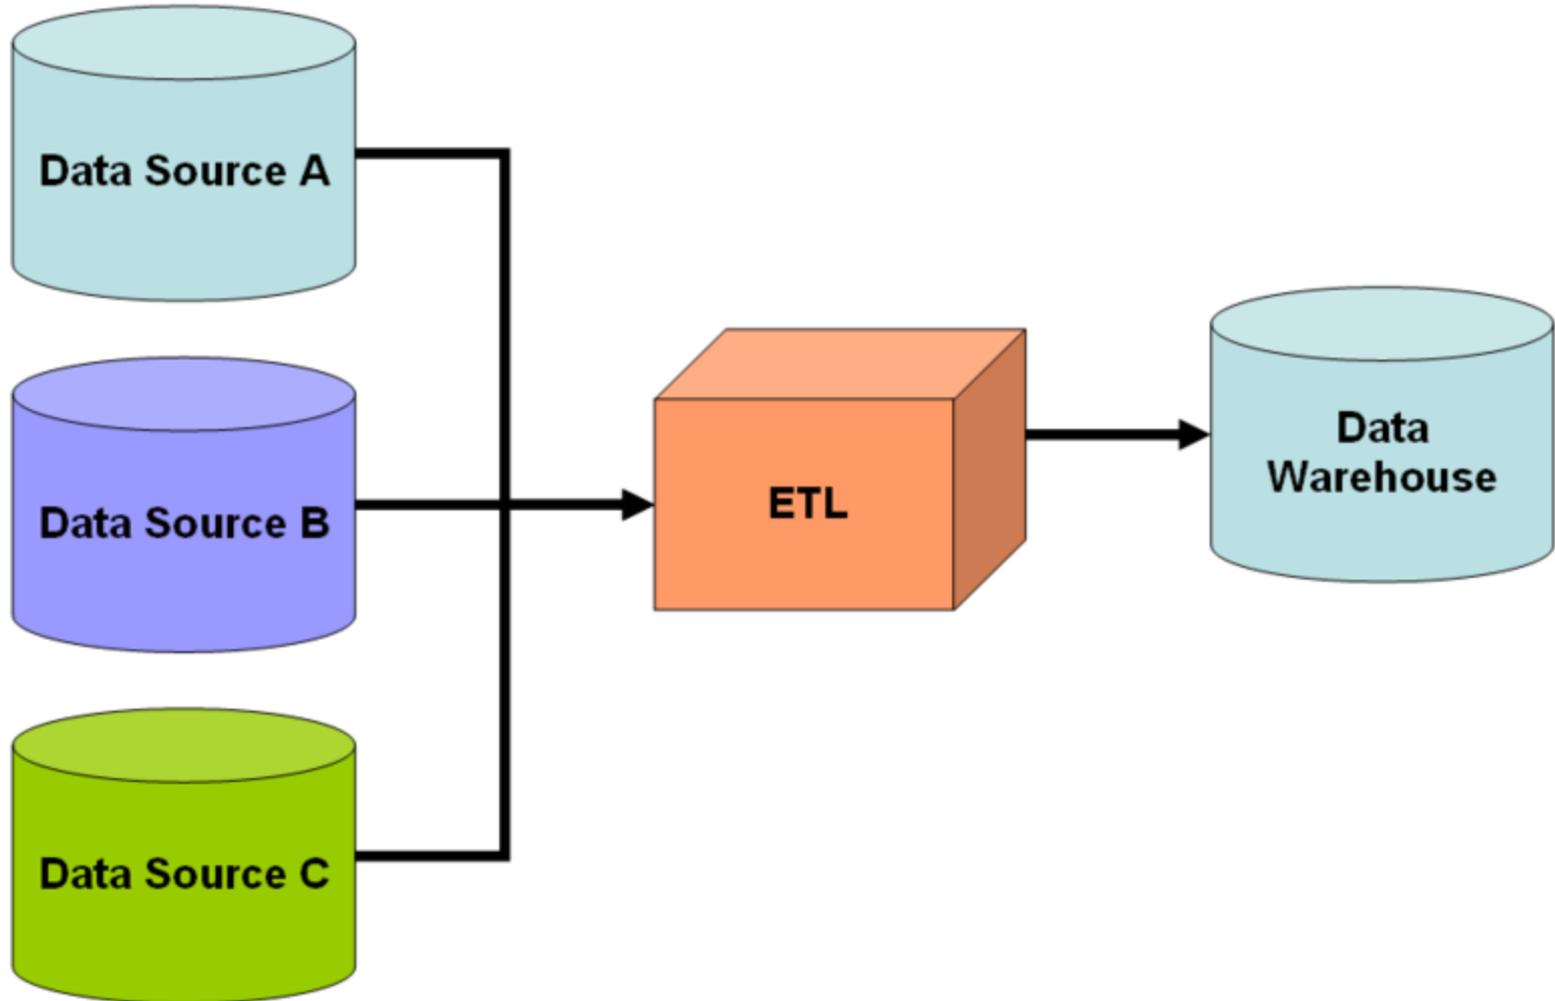
Data Integration

- Combining heterogeneous data sources under a single query interface
 - Commercial use case: two similar companies need to merge their databases
 - Scientific use case: combining research results from different bioinformatics repositories
- Types of data integration
 - Virtual integration
 - Materialized integration (Data warehousing)

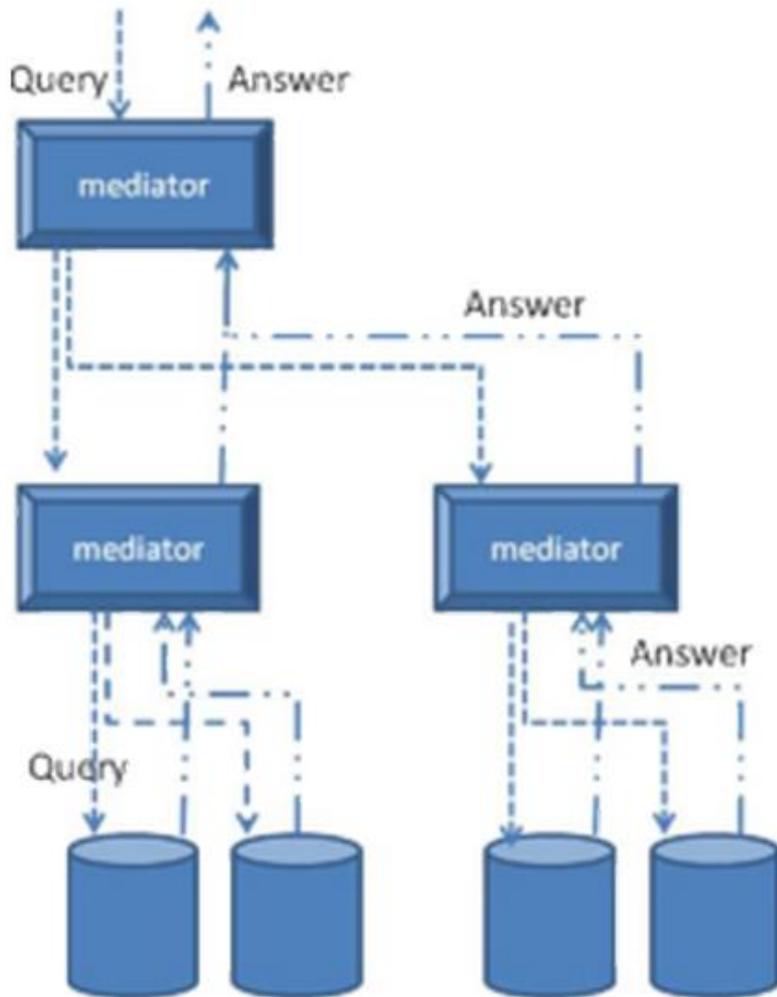
Virtual Data Integration



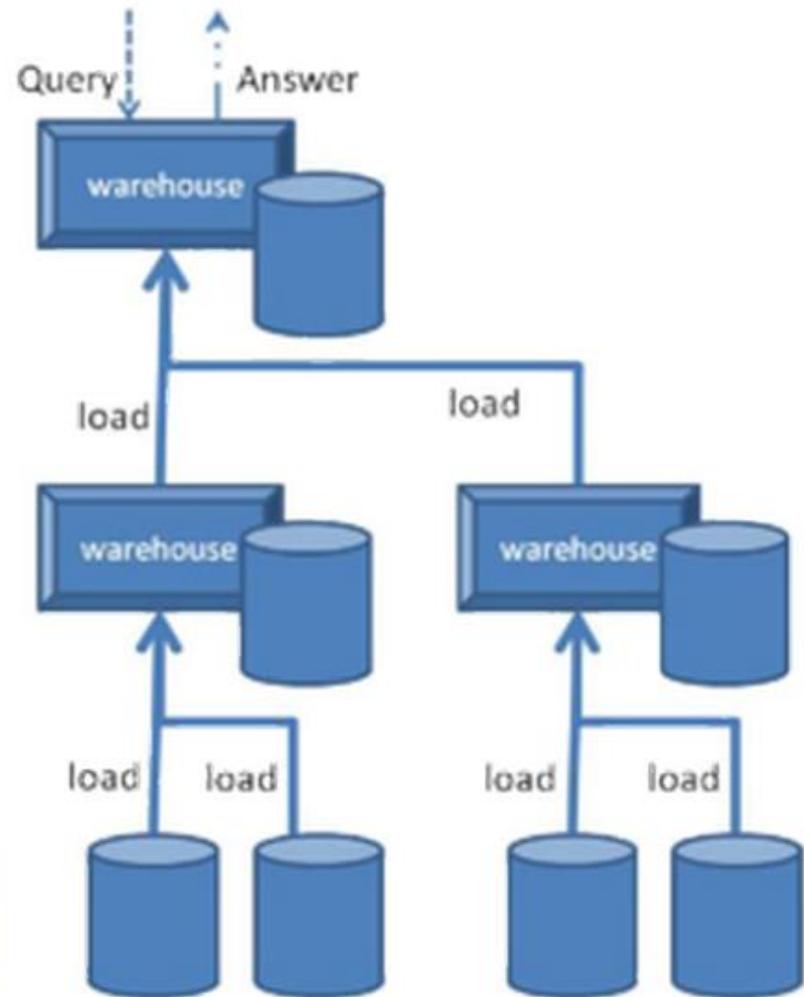
Data Warehousing (I)



Virtual versus Materialized data integration



Mediation



Warehousing

Data Warehousing (II)

- **Data warehouse** – a database that consolidates data from multiple sources
- Each resource may have a DB schema that differs from the warehouse schema. So data has to be reshaped into common warehouse schema
- Extract-Transform-Load (ETL) tools
 - cleansing operations
 - reshaping operations

What is Business Intelligence (BI)?

❑ From Encyclopedia of Database Systems:

- “[BI] refers to a set of tools and techniques that enable a company
 - to transform its business data into timely and accurate information for the decisional process,
 - to be made available to the right persons in the most suitable form.”

What is Business Intelligence (BI)?

- BI is different from Artificial Intelligence (AI)
 - AI systems make decisions for the users
 - BI systems help the users make the right decisions, based on available data
- Combination of technologies
 - Data Warehousing (DW)
 - On-Line Analytical Processing (OLAP)
 - Data Mining (DM)
 -

Why is BI Important?

- Worldwide BI revenue in 2010 = US\$ 10.5 billion (Gartner),
in 2017 = US\$ 18.3 billion (Gartner),
in 2024 = US\$ 27.36 billion (statista.com)
 - 6-7% growth each year
 - A market where players like IBM, Microsoft, Oracle, and SAP compete and invest
- BI is not only for large enterprises
 - Small and medium-sized companies can also benefit from BI
- The financial crisis has increased the focus on BI
 - You cannot afford *not* to use the “gold” in your data

Data Warehouse

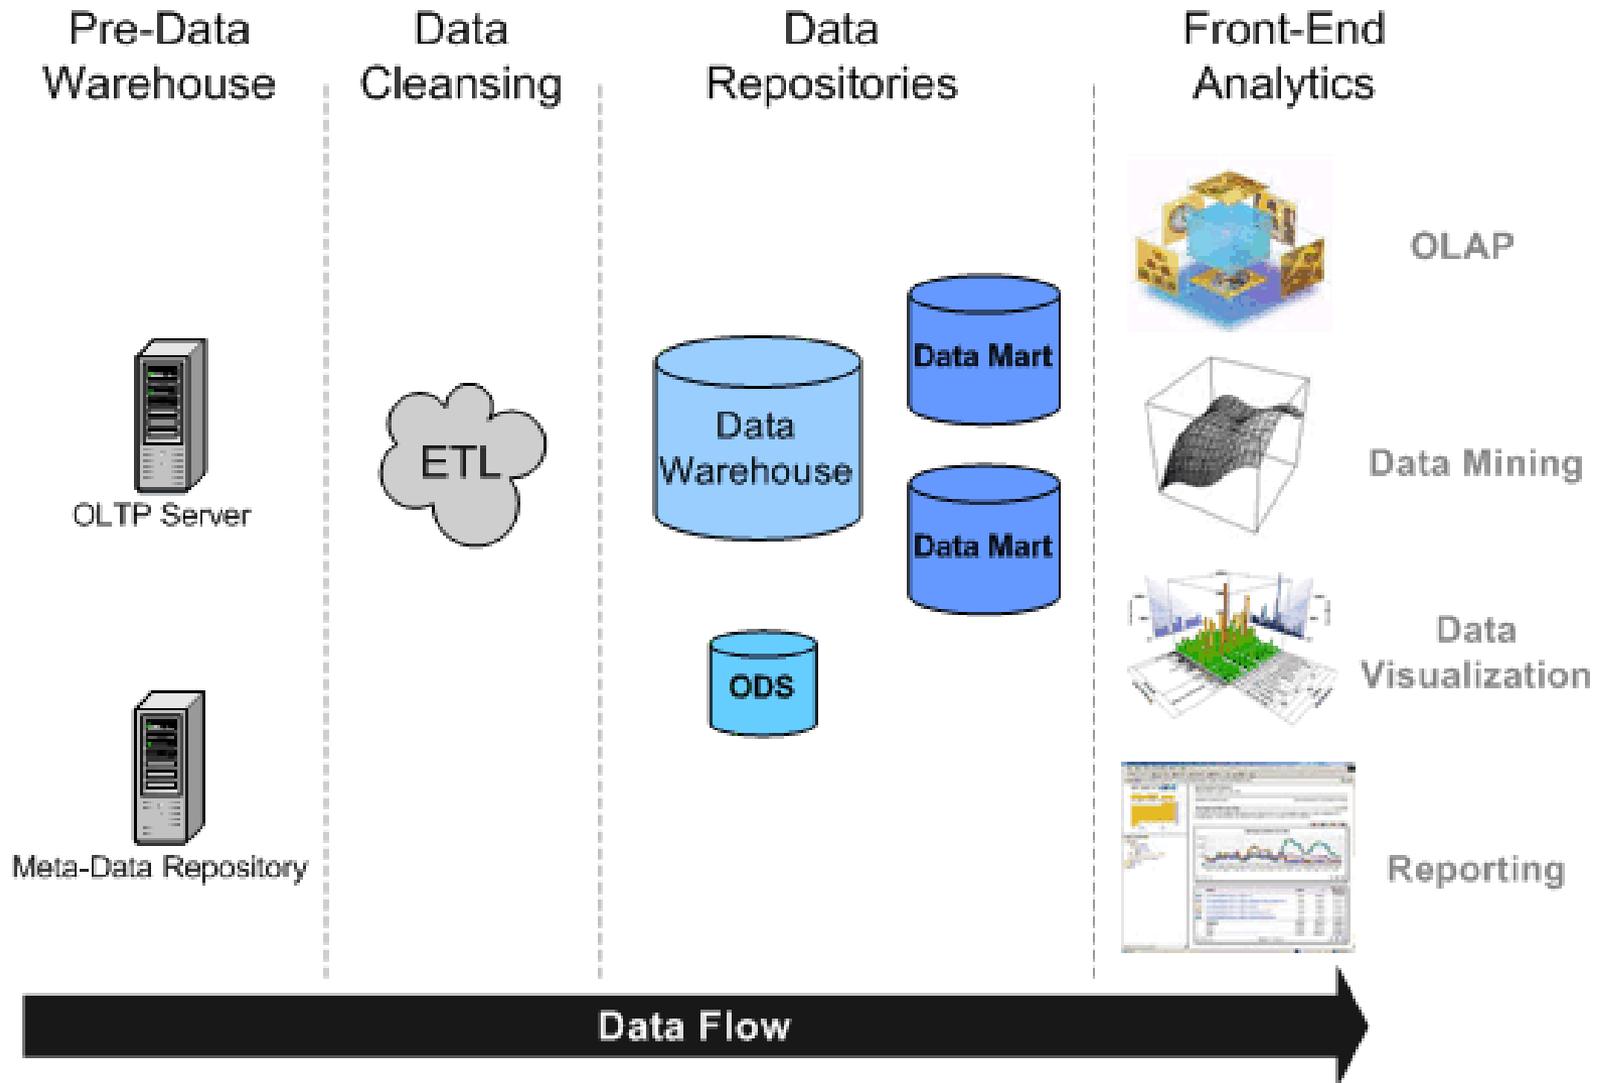
- Basically a very large database...
 - Not all very large databases are DW, but all data warehouses are pretty large databases
 - Nowadays a warehouse is considered to start at around a TB and goes up to several PB
 - It spans over several servers and needs an impressive amount of computing power
- More specific, a collective data repository
 - Containing snapshots of the operational data (history)
 - Obtained through data cleansing (Extract-Transform-Load process)
 - Useful for analytics



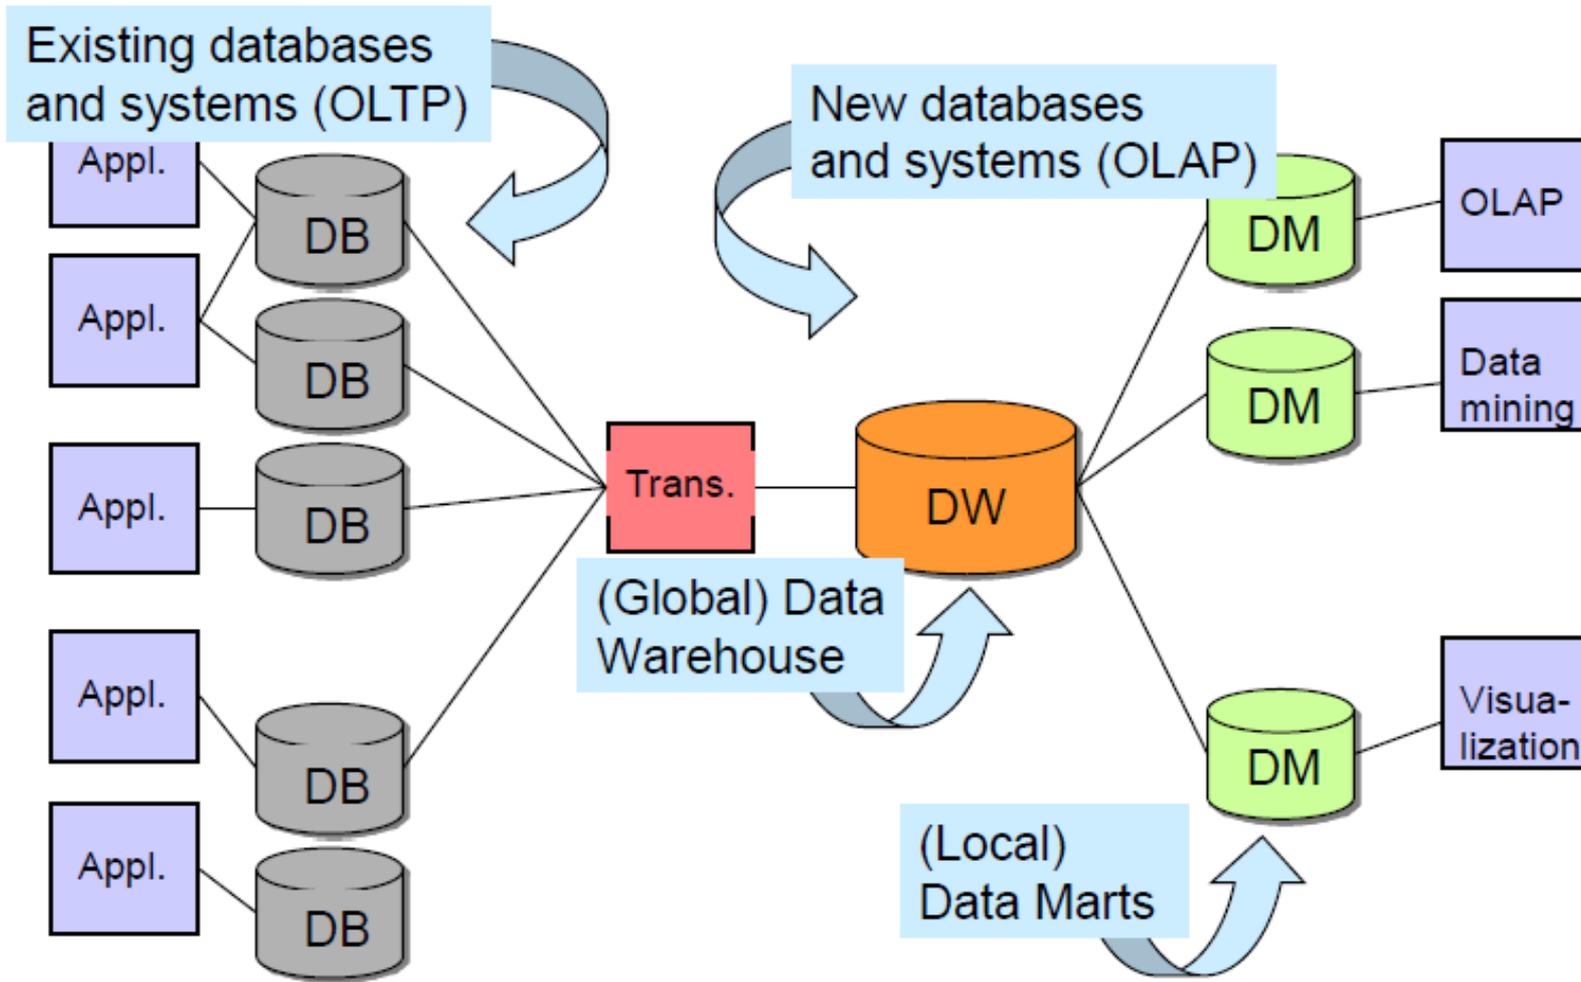
DW: Purpose and Definition

- DW is a **store of information** organized in a unified data model
- Data collected from a number of different sources
 - Finance, billing, website logs, personnel, ...
- Purpose of a data warehouse (DW): support **decision making**
- Easy to perform advanced analysis
 - Ad-hoc analysis and reports
 - ◆ We will cover this soon
 - Data mining: discovery of hidden patterns and trends

Data Warehouse: Data Flow



DW Architecture – Data as Materialized Views



Analogy: (data) producers ↔ warehouse ↔ (data) consumers

ETL (I)

- Short for Extract, Transform and Load
- Database functions that are combined into one tool to pull data out of productive databases and place it into the DW
 - Migrate data from one database to another, to form data marts and data warehouses
 - Convert databases from one format or type to another
- ETL is used to integrate heterogeneous systems
 - With different DBMS, operating system, hardware, communication protocols
- ETL challenges
 - Getting the data from the source to target as fast as possible
 - Allow recovery from failure without restarting the whole process

ETL (II)

- “Getting multidimensional data into the DW”
- Problems
 1. Data from different sources
 2. Data with different formats
 3. Handling of missing data and erroneous data
 4. Query performance of DW
- ETL
 - Extract (for problem #1)
 - Transformations / cleansing (for problems #2, #3)
 - Load (for problem #4)
- The most time-consuming process in DW development
 - 80% of development time spent on ETL

OLAP

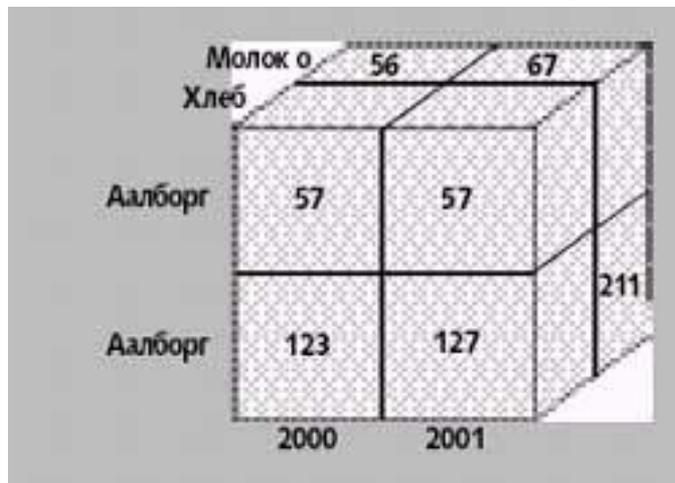
- OLTP (OnLine Transaction Processing)
 - Typically for data entry / retrieval and transaction processing
 - Works on the operational data stores (ODS) and represents day-to-day operational business activities
 - Purchasing, sales, production distribution, ...
 - Reflects only the current state of the data
- OLAP (OnLine Analytical Processing)
 - Provides information for activities like
 - Enterprise resource planning, capital budgeting, marketing initiatives,...
 - Represents front-end analytics based on a DW repository
 - Used for reporting and it is decision oriented
 - Needs to provide timely, accurate and understandable information
 - minutes, hours or even longer

OLTP vs. OLAP

	OLTP	OLAP
Target	operational needs	business analysis
Data	small, operational data	large, historical data
Model	normalized	denormalized/ multidimensional
Query language	SQL	not unified – but MDX is used by many
Queries	small	large
Updates	frequent and small	infrequent and batch
Transactional recovery	necessary	not necessary
Optimized for	update operations	query operations

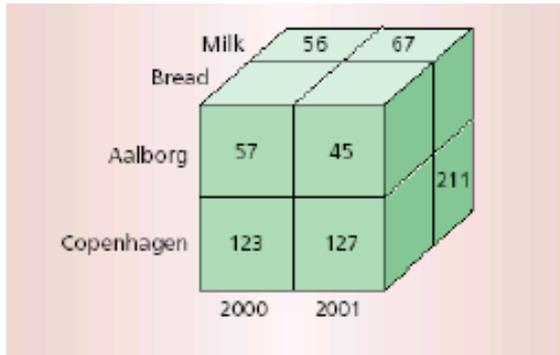
OLAP Data Cubes

- Факт – набор значений измерений (dimensions), которому сопоставлен набор параметров (measure)
 - (покупка — факт, объем покупки и стоимость — параметры, тип продукта, время и место покупки — измерения)
- Запросы агрегируют значения параметров по всему диапазону измерения (общий месячный объем продаж данного продукта)
- Измерения организуются в иерархию, состоящую из нескольких уровней, каждый из которых представляет уровень детализации, требуемый для соответствующего анализа

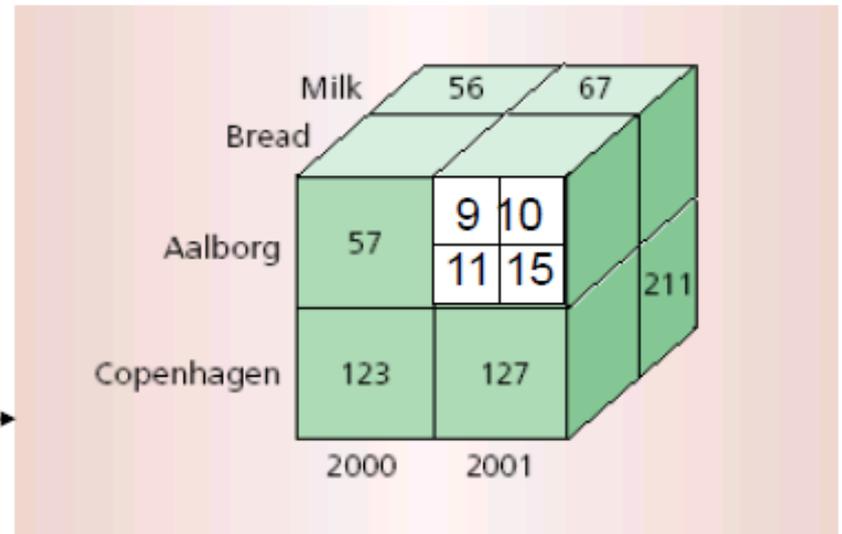
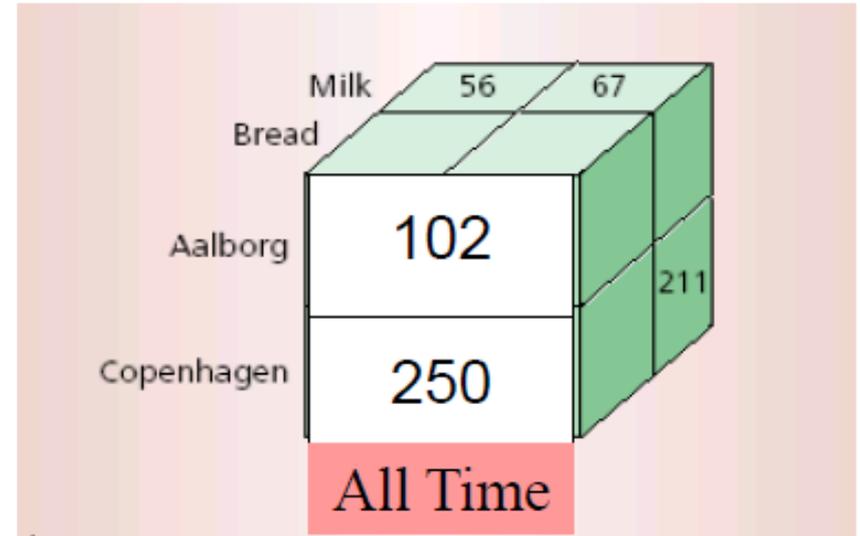


- Реализации: Cognos Powerplay, Oracle Database OLAP Option, Microsoft Analysis Services, etc.

On-Line Analytical Processing (OLAP)



- On-Line Analytical Processing
 - Interactive analysis
 - Explorative discovery
 - Fast response times required
- OLAP operations/queries
 - Aggregation, e.g., SUM
 - Starting level, (Year, City)
 - ◆ Roll Up: Less detail
 - ◆ Drill Down: More detail
 - Slice/Dice: Selection, Year=2000



Data Warehouse Appliances

- Integrated set of
 - Servers
 - Storage
 - operating system(s)
 - DBMS and software
- specifically pre-installed and pre-optimized for data warehousing (DW)

- ❑ Use massively parallel processing (MPP) architectures to provide high query performance and platform scalability
 - ❑ independent processors or servers executing in parallel
 - ❑ "shared-nothing architecture" where each server operates self-sufficiently and controls its own memory and disk

What's MPP in Data Warehousing ?

MPP (massively parallel processing) data warehouse systems are different from SMP (symmetric multiprocessing) databases:

1. Shared-nothing architectures, with **no single point of failure** and often hot-swappable components
2. **Scale horizontally** by adding nodes, rather than moving to a server with more CPUs or higher storage capacity
3. Breaks a **large queries** across nodes for simultaneous processing
4. Capable of higher data **ingestion rates** through parallelized data movement

Comparing Architectures

	Teradata	Netezza	Greenplum	Oracle Exadata
Hardware	Custom MPP, Shared Nothing	Custom MPP: SPU + FPGA logic	Commodity hardware	Custom Hybrid MPP, Shared Everything
Type of processing	OLTP or OLAP, Can handle high user load	OLAP, Assumes few users for heavy analytics	OLAP	OLTP or OLAP, Ideal for Oracle shops
Inception / Maturity	1979	2000	2003	2008
Performance & maintenance	Auto-recommended optimization, columnar compression available	No need for performance tuning, Must manually reclaim space	Based on PostgreSQL, but optimized for MPP and enterprise maint.	Some tuning required, Indexing for OLTP

Shared nothing VS Shared Everything

- In a *shared nothing* design, each node of an MPP cluster owns a set of data and handles all processing of data that resides on that node
- In a *shared everything* design, each node of the cluster can access all data, and data can be processed on any node
- *Exadata* is a hybrid of both approaches:
 - Each Exadata storage cell owns a set of data and handles all processing on that data
 - Database functions that are more data intensive (process large volumes of data) are offloaded into the storage layer
 - Other functions of the database are processed in the central core or processing cluster of Exadata
- <https://chriscraftoracle.wordpress.com/2018/02/09/what-makes-exadata-faster/>

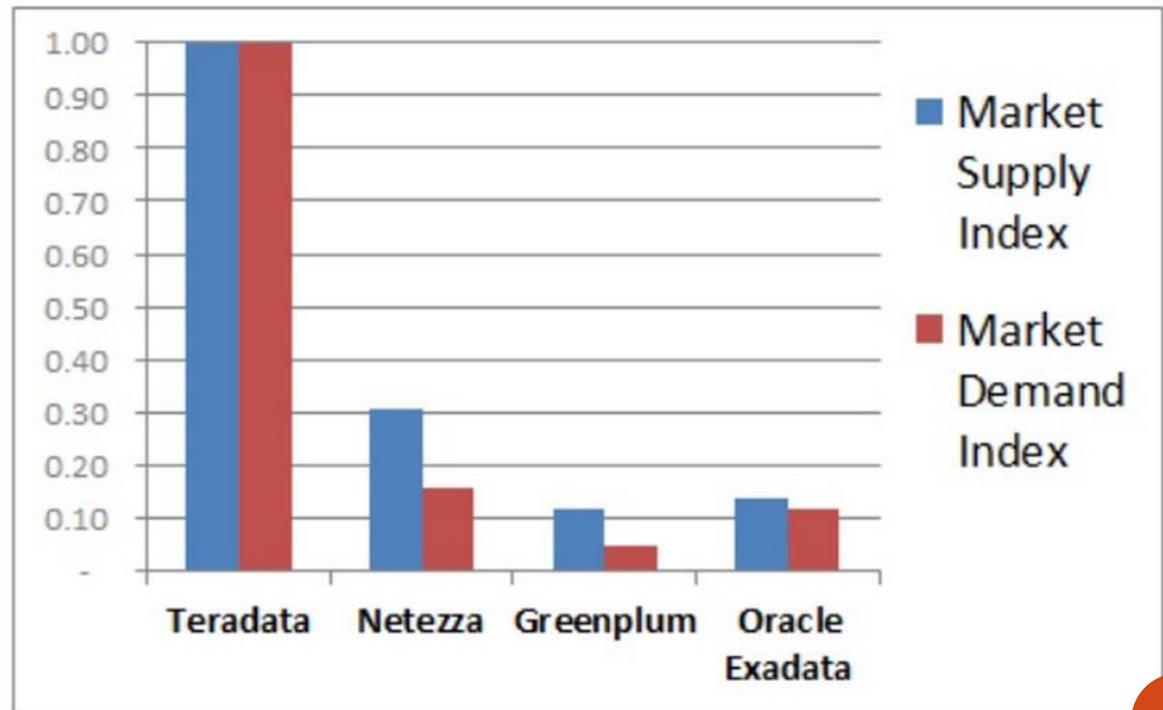
Comparing Costs

Teradata	<p>Hardware and licenses the most expensive of all options. Staff costs can be expensive and it takes a great deal of effort to configure and administer.</p>	
IBM Netezza	<p>Hardware and licenses used to be much less than Teradata, but prices have been converging. Some of the highest staff cost due to scarcity, but that's tempered by lower effort for configuration and admin of single purpose appliance.</p>	
Greenplum	<p>Commodity hardware. Moderately priced licenses. Few Greenplum specialists, but can be staffed by PostgreSQL DBAs and devs.</p>	
Oracle Exadata	<p>Moderately expensive hardware and license, compared to the shared nothing options. Lots of Oracle expertise in the market makes staffing costs moderate.</p>	
Hadoop HBase	<p>Commodity hardware and no license cost, resulting in lowest up-front cost. But requires highly technical staff and implementation is less productive than more mature options.</p>	

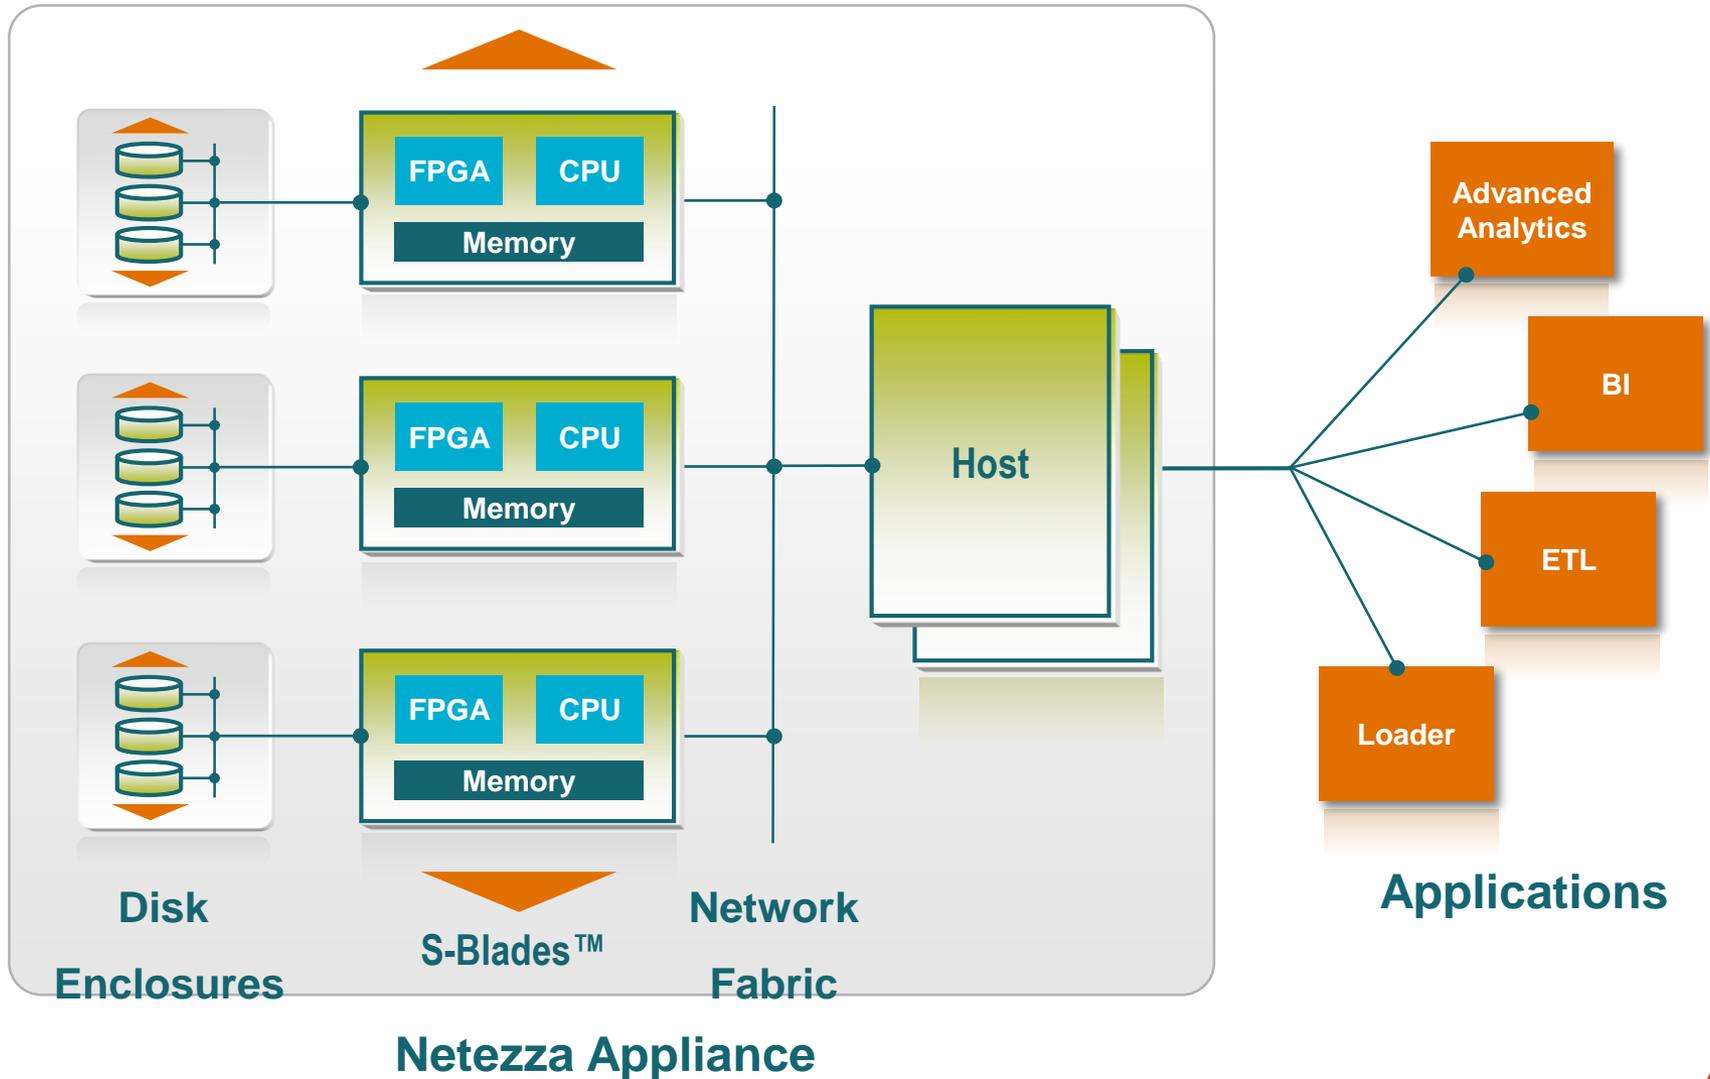
Popularity

Market Comparison

Teradata has been around for many years longer than the alternatives and still dominates the market in terms of install base (3 times next rival) and vibrant development community (6 times next rival).



Netezza AMPP архитектура



Спецификация TwinFin 24



■ 16 (8*2) Disk Enclosures:

- 192 (96*2) 1TB SAS Drives
- (8 hot spares)
- RAID 1 Mirroring

■ 2 Hosts (Active-Passive):

- 24 Cores (Quad-Core Intel 2.6 GHz)
- 96 GB Memory
- 4x146 GB SAS Drives
- Red Hat Linux 5 64-bit
- 10G Internal Network

■ 24 Netezza S-Blades:

- 192 Core's (Intel Quad-Core 2.5 GHz)
- 192 FPGA's (125 MHz)
- 384 GB DDR2 RAM (1+TB compressed)
- Linux 64-bit Kernel

- User Data Capacity: 250 TB
- Data Scan Speed: 290 TB/hr
- Load Speed (per system): 2.0 TB/hr

- Power/Rack: 7,400 Watts
- Cooling/Rack: 25,500 BTU/Hour