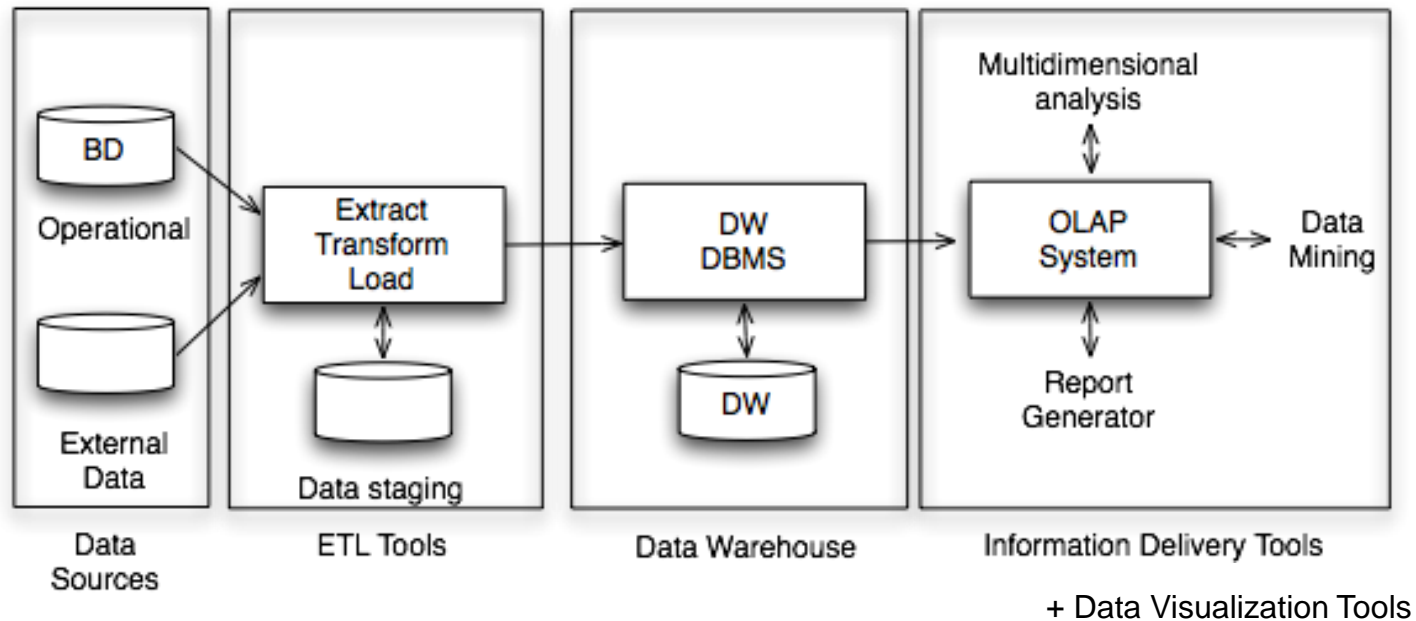


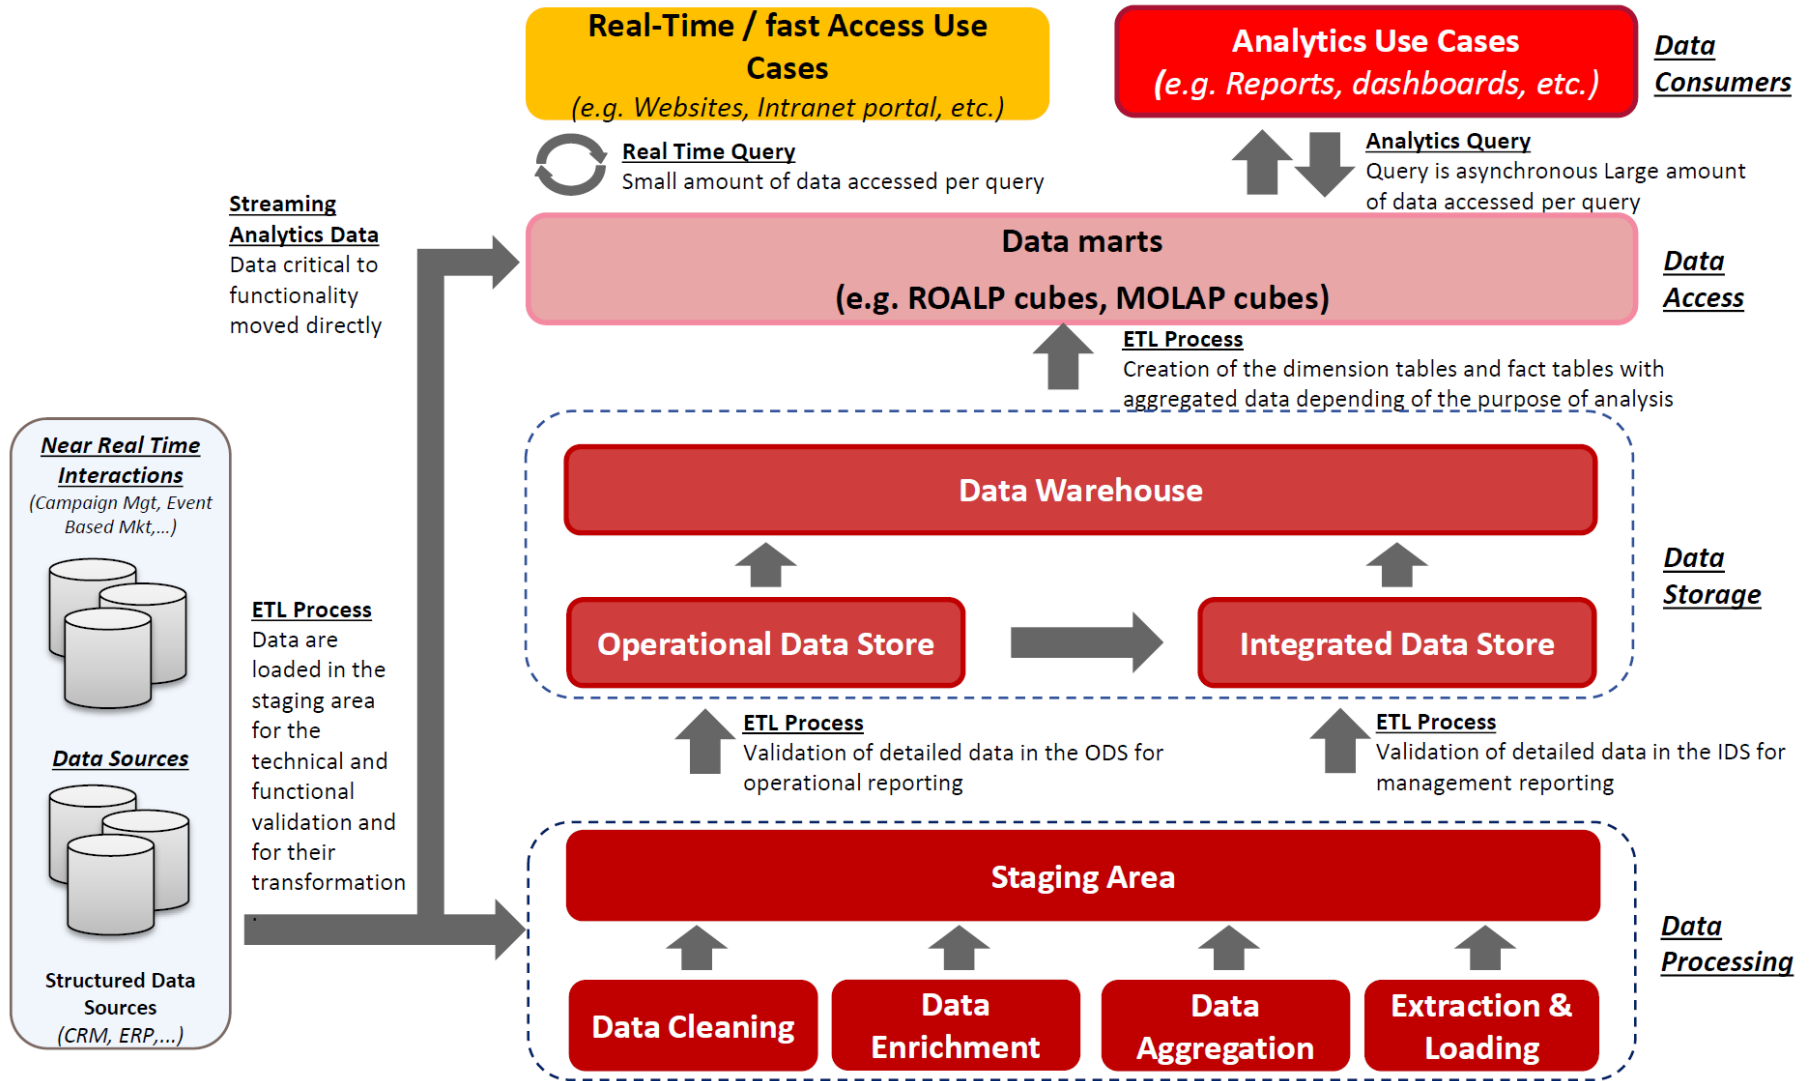
Decision support databases: Trends in Data Warehousing



Traditional DSS Architecture




Traditional DSS Architecture



Operational Data Store

An ODS is an architectural construct containing subject oriented, integrated, **volatile**, **current** (or near-current), and **detailed** operational information.

It is used for operational reporting, controls and decision making. Require **faster data availability** than in the DW (Zero-Latency Enterprise).



Operational Data Store

An ODS is an architectural construct containing subject oriented, integrated, **volatile**, **current** (or near-current), and **detailed** operational information.

It is used for operational reporting, controls and decision making. Require faster data availability than in the DW (Zero-Latency Enterprise).

Retail

- ▶ What inventory items should I be adjusting throughout the day?
- ▶ How can my customers track their own orders through the Web?
- ▶ What are my customers ordering across all subsidiaries?
- ▶ What is the buying potential of my customer at the point of sale?

Volume and Velocity

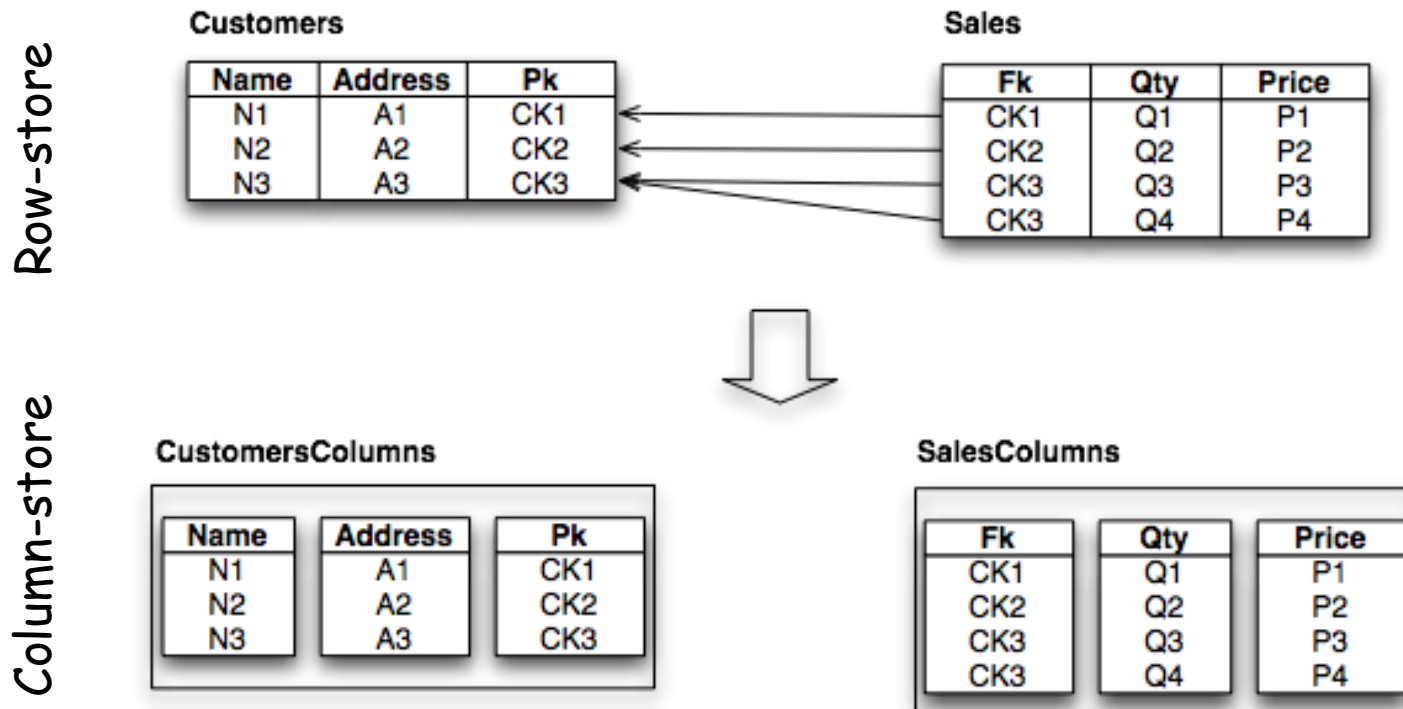
Scalability to large volume and to real-time updates of facts:

- Column-store DWs (Volume)
- In-memory DWs (Velocity)

Column-oriented DWs

Change a fundamental assumption of traditional DBMS:

Vertical table partitioning: store data by columns, not by rows



Commercial: Oracle, SQL Server, Sybase IQ, HP Vertica, SAP Hana, SADAS

Column-oriented DWs

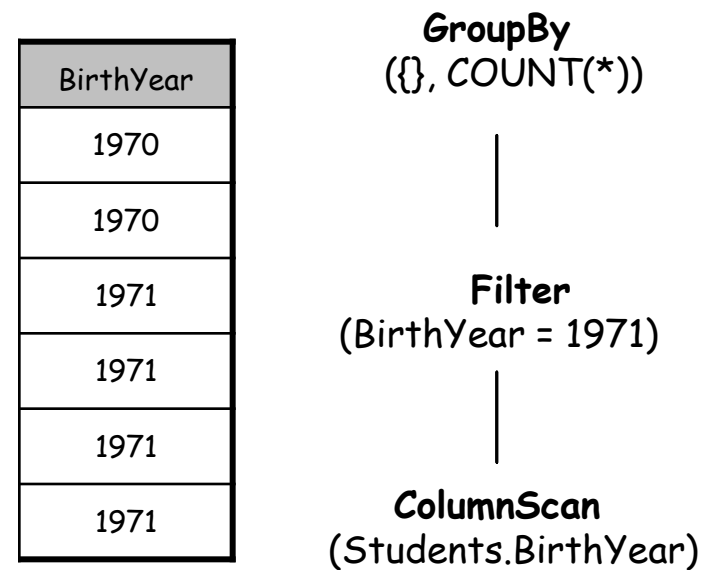
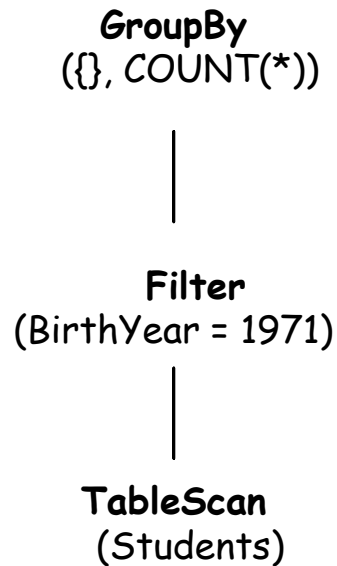
RID	StudCode	City	BirthYear	BirthYear	BirthYear	N	City	ID	City
1	100	MI	1970	1970	1970	2	2	1	FI
2	101	PI	1970	1970	1971	4	3	2	MI
3	102	PI	1971	1971	1971		3	3	PI
4	104	FI	1971	1971	1971		1		
5	106	MI	1971	1971	1971		2		
6	107	PI	1971	1971	1971		3		

Space gains: By compressing each column using a compression method that is most effective for it, *substantial reductions in the total size* of data on disk can be achieved.

Column-oriented DWs

```
SELECT COUNT(*)  
FROM Students  
WHERE BirthYear = 1971
```

RID	StudCode	City	BirthYear
1	100	MI	1970
2	101	PI	1970
3	102	PI	1971
4	104	FI	1971
5	106	MI	1971
6	107	PI	1971

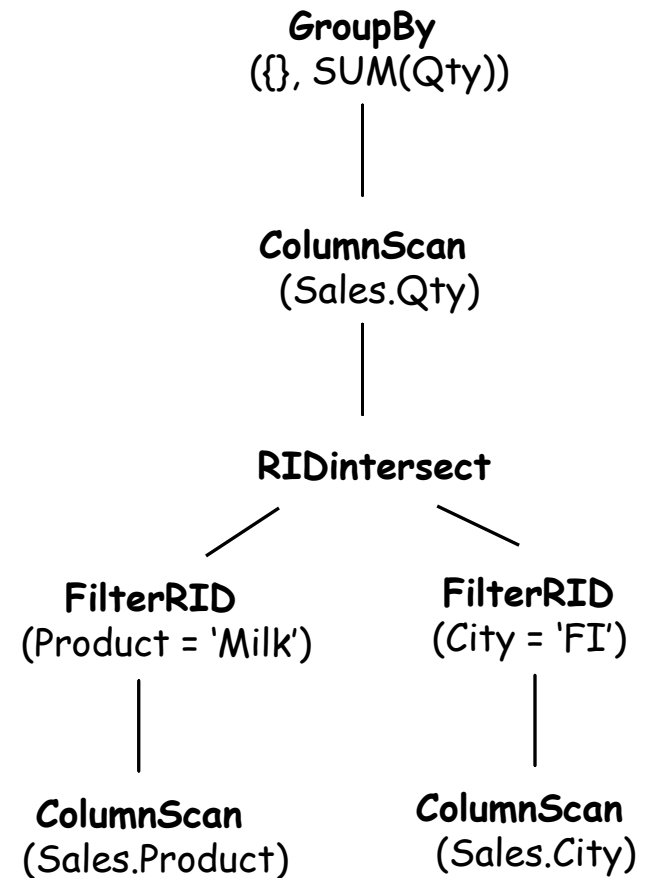


Efficiency gains: By scanning only the required columns, *substantial reductions in the total time* of accessing data on disk can be achieved.

Column-oriented DWs

```
SELECT SUM(Qty)
FROM Sales
WHERE Product = 'Milk' AND City = 'FI'
```

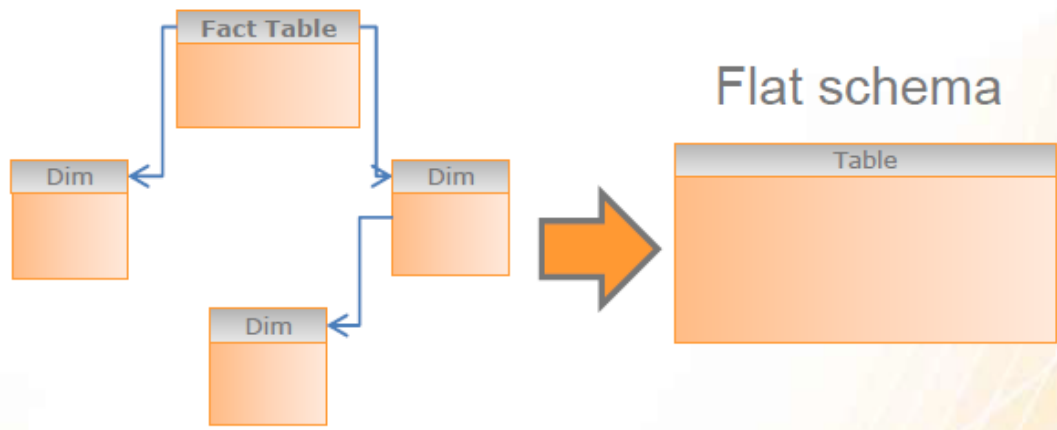
RID	Product	City	Qty
1	Milk	MI	3
2	Bread	PI	1
3	Meat	PI	2
4	Milk	FI	2
5	Milk	MI	1
6	Bread	PI	1



Late materialization: operators work mostly on columns, using RIDs as input/output to other operators

Column-oriented DWs

Denormalization: fact table includes all dimensional attributes + BM index on all attributes.



RID	StudCode	City	BirthYear
1	100	MI	1970
2	101	PI	1970
3	102	PI	1971
4	104	FI	1971
5	106	MI	1971
6	107	PI	1971

City
MI
PI
PI
FI
MI
PI

City	BM
FI	000100
MI	100010
PI	011001

In-memory DWs

For the last 30 years main memory prices have dropped by a factor of 10 every 5 years. Moreover, ends of Moore's law boosted multi-core processors.



2TB Kit (8 x 256GB) DDR4-3200 PC4-25600 ECC Load Reduced Memory for ASRock Rack ROMED8-2T AMD EPYC Board by NEMIX RAM

Brand: NEMIXRAM

Price: **\$20,792.99** + \$41.00 shipping

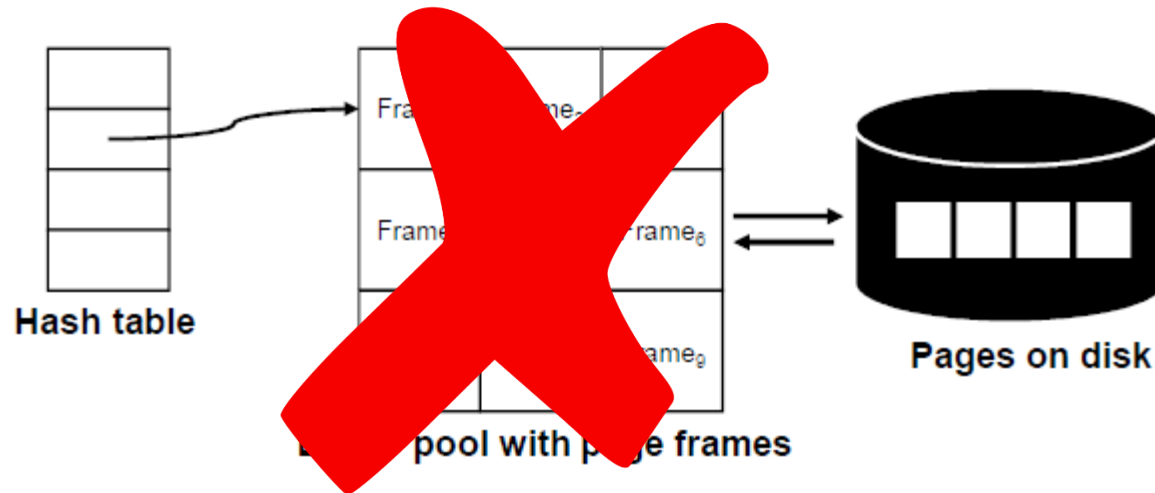
- Compatible with ASRock Rack ROMED8-2T AMD EPYC Board
- 2TB (8 x 256GB)
- DDR4 3200 (PC4 25600)
- ECC Load Reduced LRDIMM / 288-Pin SDRAM 1.2V
- Lifetime Replacement Warranty

DDR4 RAM throughput approx 50 GB/s	- highly expensive
SSD throughput approx 2 GB/s	- expensive
SATA HD throughput approx 0.5 GB/s	- cheap

In-memory DWs

Change another fundamental assumption of traditional DBMS:

In-memory DWs: store all data in main memory, no I/O transfer from disk.




It also requires new query optimization (compilation to sw code).

Applications: real-time analytics (e.g., stock exchange, sensor monitoring, telecom industry, etc.)

Commercial: Oracle, SQL Server, SAP Hana,

Query compilation example

```
SELECT *  
FROM R, S, T  
WHERE T.x=7 and S.y=3 and R.z>5 and T.B=S.B and S.A=R.A
```



```
initialize memory of hash tables  $\bowtie_A, \bowtie_B$   
for each tuple  $t$  in  $T$   
  if  $t.x = 7$   
    materialize  $t$  in hash table of  $\bowtie_B$   
for each tuple  $s$  in  $S$   
  if  $s.y = 3$   
    materialize  $s$  in hash table of  $\bowtie_A$   
for each tuple  $r$  in  $R$   
  if  $r.z > 5$   
    for each match  $s$  in  $\bowtie_B [r.B]$   
      for each match  $t$  in  $\bowtie_A [s.A]$   
        output  $r \circ s \circ t$ 
```

Figure 4.16: Generated pseudo-code

Volume and Velocity

Scalability to large volume and to real-time updates of facts:

- Column-store DWs (Volume)
- In-memory DWs (Velocity)
- DW Appliances/in the Cloud (Volume/Velocity)

DW Appliance

A **DW appliance** includes an integrated set of servers, storage, operating systems, and databases, with pre-configuration optimized for performance

- solve the complexity of assembling HW, OS, DBMS, and tools for ETL, OLAP, and reporting

*Nodes	Eight MPP nodes per cabinet with Intel Westmere Dual Six Core Xeon Processors
Storage	(192) 300GB, 600GB or 900GB SAS drives per Cabinet
Total User Data Capacity	18.2TB Per Cabinet – 300GB 36.5TB Per Cabinet – 600GB 54.9TB Per Cabinet – 900GB (uncompressed)
Scalability	Scales up to 46 nodes: • 105TB (300GB) • 210TB (600GB) • 315TB (900GB)*
Availability	RAID 1 Disk Mirroring, Node Failover with Cliquing, BAR
Operating System	SUSE Linux
System Management	Single Operational View Across Complete System via Teradata Viewpoint
Interconnect	Teradata BYNET®



KEY FEATURES

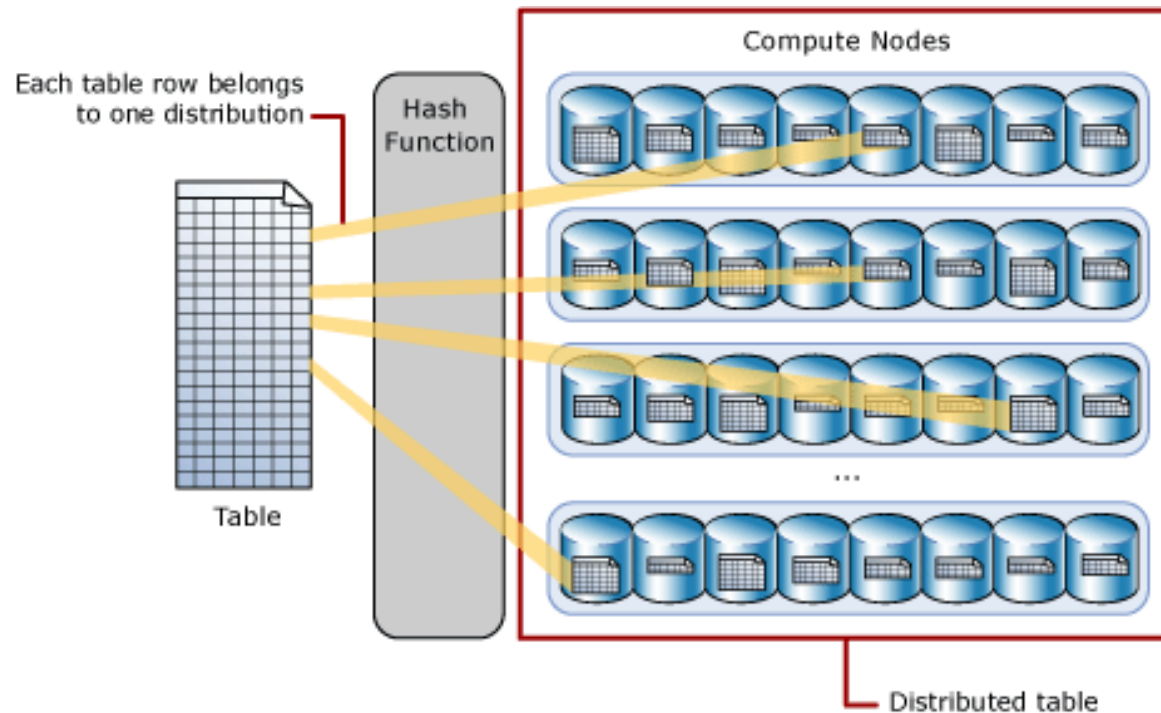
- Up to 684 CPU cores and 14.6.TB memory per rack for database processing
- Up to 288 CPU cores per rack dedicated to SQL processing in storage
- From 2 to 19 database servers per rack
- From 3 to 18 Oracle Exadata Storage Servers per rack
- Up to 230 TB of Flash Storage per rack
- 40 Gb/second (QDR) InfiniBand Network
- Uncompressed and mirrored usable capacity of up to 385 TB per rack
- Hybrid Columnar Compression often delivers 10X-15X compression ratios
- Complete redundancy for high availability
- Oracle Linux

Commercial: IBM Netezza, Teradata, Oracle Exadata, ...

Massively Parallel Processing

Distribution of data (e.g., Hash distribution of records) and distributed execution of queries

- Major DBMS (Oracle, SQL Server, DB2) offer MPP storage and distributed queries



DW in the Cloud

Free from managing physical data centers (managed DBMS), elastic to business changes, MPP, Columnar DB, etc.

- SaaS model (pay per use), Lock-in effect
- [Amazon Redshift](#), [Microsoft Azure Synapse](#), [Google BigQuery](#), [Snowflake Data Cloud](#)

<p>Business Intelligence & Machine Learning</p>			<p>AWS Marketplace 250+ solutions</p>
<p>Databases</p>	<p>Analytics</p>	<p>Blockchain</p>	<p>730+ Database solutions</p> <p>600+ Analytics solutions</p> <p>25+ Blockchain solutions</p>
<p>Data Lake</p>		<p>20+ Data lake solutions</p>	
<p>Data Movement</p> <p>Database Migration Service Snowball Snowmobile Kinesis Data Firehose Kinesis Data Streams Data Pipeline Direct Connect</p>			<p>30+ solutions</p>

Volume and Velocity

Scalability to large volume and to real-time updates of facts:

- Column-store DWs (Volume)
- In-memory DWs (Velocity)
- DW Appliances/in the Cloud (Volume/Velocity)
- Complex event processing (Velocity)

Complex Event Processing

Real-Time / fast Access Use Cases

(e.g. Websites, Intranet portal, etc.)



Real Time Query

Small amount of data accessed per query

Data marts

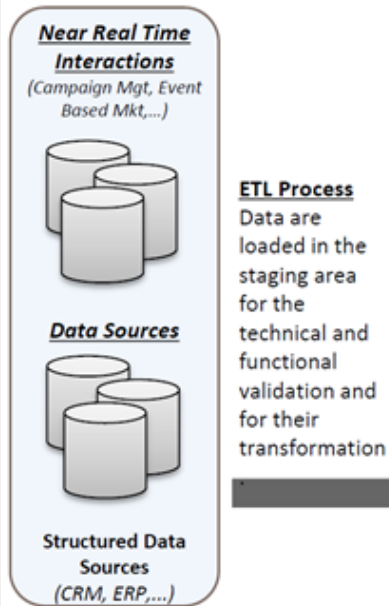
(e.g. ROALP cubes, MOLAP cubes)

Streaming Analytics Data
Data critical to functionality moved directly

CEP tracks and analyse streams of events collected from operational sources with the purpose of detecting patterns and raising alarms in real time.

Application areas: stock market analysis, mobility analysis, production monitoring, credit card fraud detection, security intrusion detection, logistics, ...

More at Business Process Modeling course, next year.



SQL triggers for event processing

Triggers (or ECA rule, or **event-condition-action** rule) let the user decide **when and what** to check for a condition.

Event: typically a type of database modification, e.g., "insert"

Condition: Any SQL boolean-valued expression.

Action: Any SQL statements.

```
CREATE TRIGGER SalesTrig
  AFTER INSERT ON Sales
  REFERENCING NEW ROW AS NewTuple
  FOR EACH ROW
  WHEN (NewTuple.ProductName NOT IN
        (SELECT name FROM Products))
  INSERT INTO Products(name)
  VALUES(NewTuple.ProductName);
```

The event

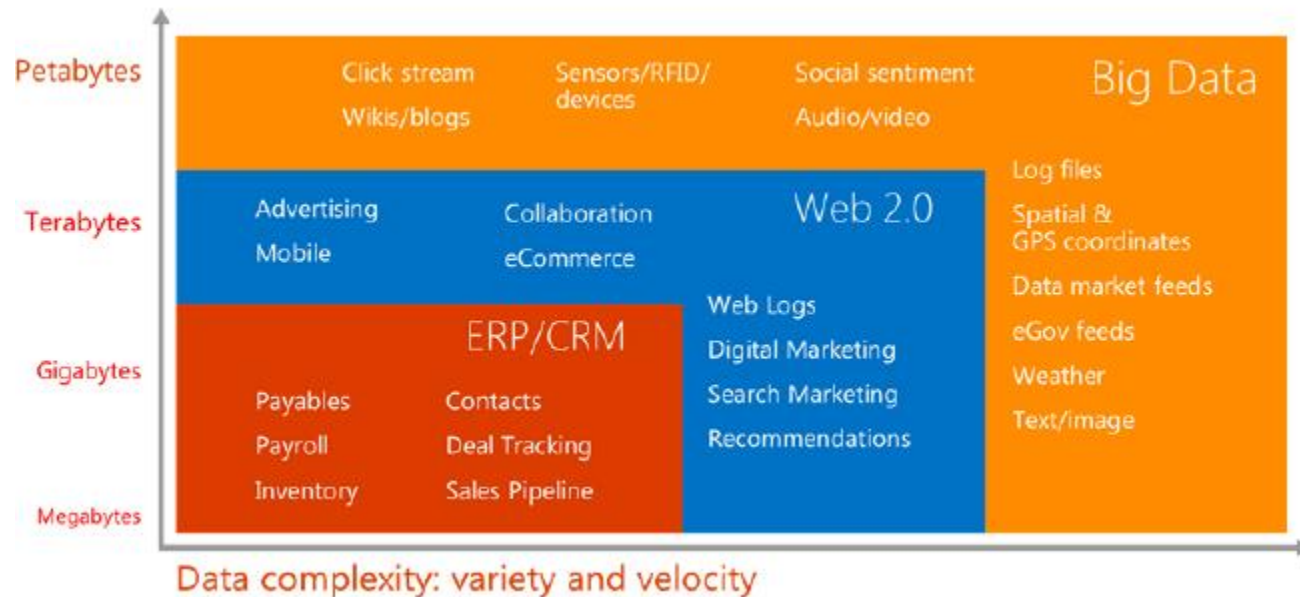
The condition

The action

Volume and Velocity and ... Variety

Variety: ability to store, process and query both structured and unstructured data types:

- structured: tuples
- unstructured: text, graphs, images, sound, videos, trajectories, ...



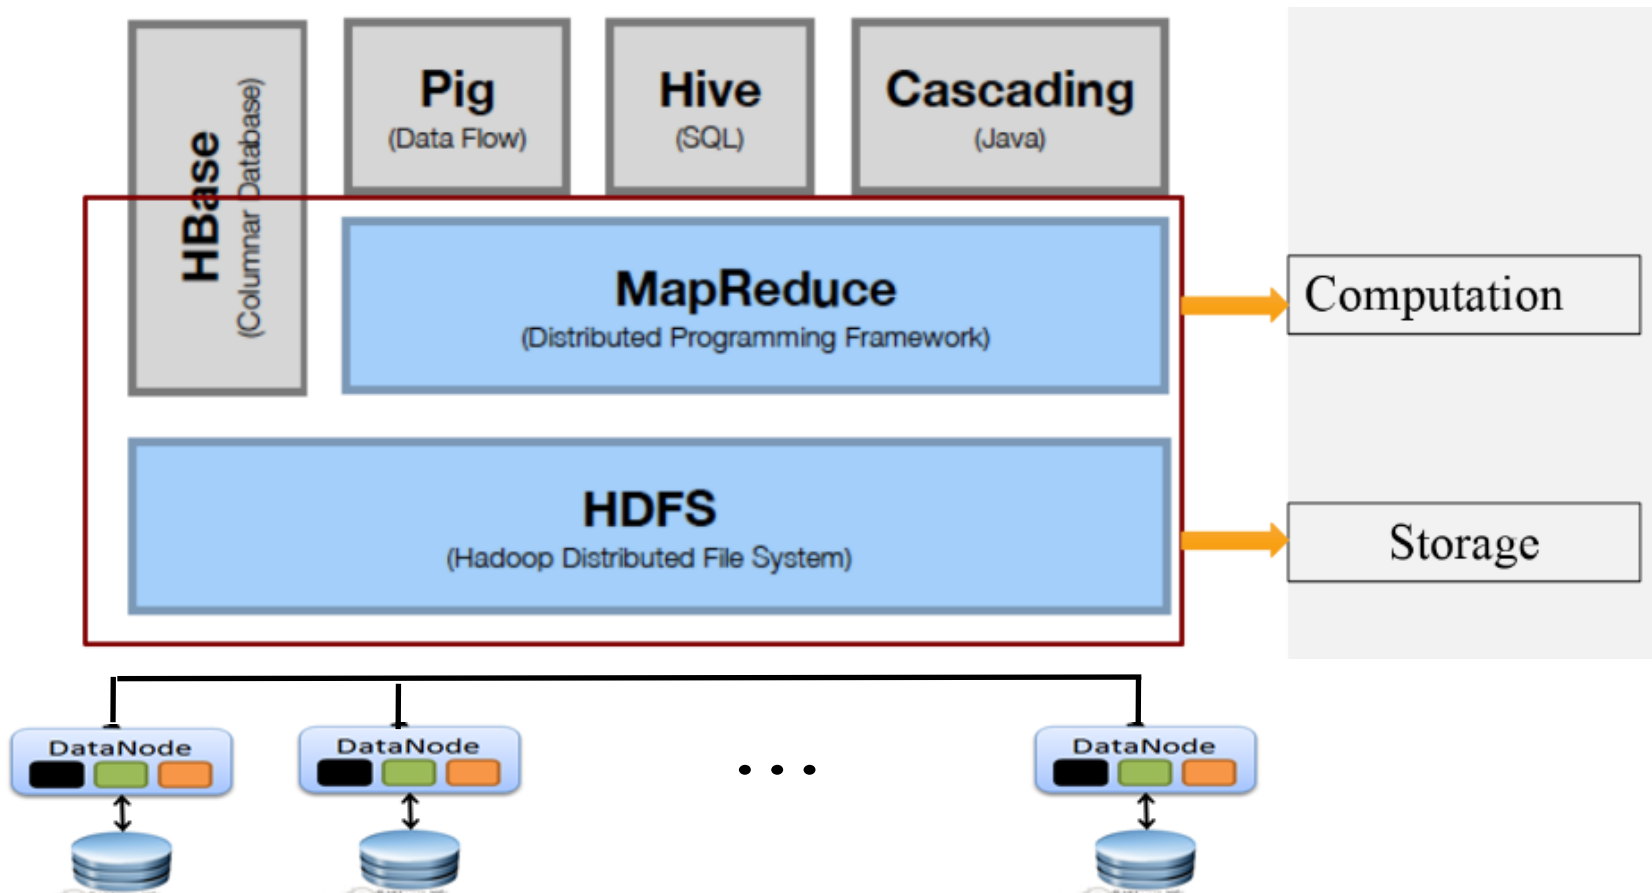
Big Data Framework

Key Differences Between the Data Lake and Data Warehouse

DATA WAREHOUSE	vs.	DATA LAKE
Structured, processed	DATA	Structured/semi-structured/unstructured/raw
Schema-on-write	PROCESSING	Schema-on-read
Expensive for large data volumes	STORAGE	Designed for low-cost storage
Less agile, fixed configuration	AGILITY	Highly agile, configure and reconfigure as needed
Mature	SECURITY	Maturing
Business pros	USERS	Data scientists et al.

Analysis Source: "A Big Data Cheat Sheet: What Marketers Want to Know" by Tamara Dull

Hadoop Ecosystem

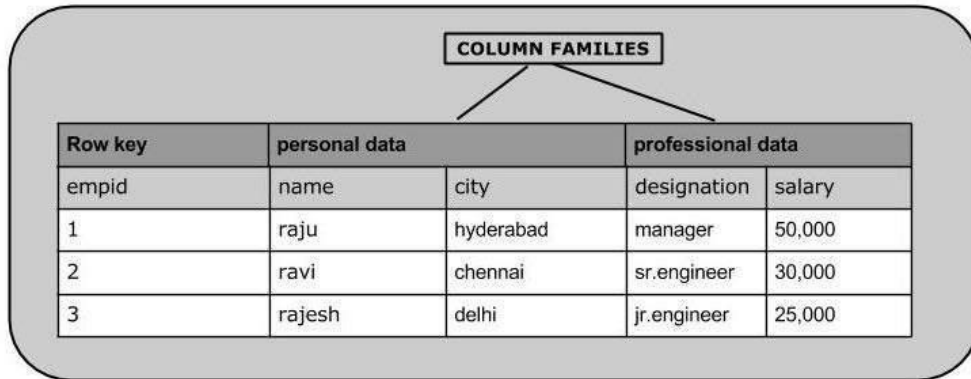


<http://hbase.apache.org>

HBase: columnar database, distributed, scalable, big data store. A NoSQL data store, i.e., not a relational data model.

NoSQL = Not Only SQL

- **Column:** Accumulo, Cassandra, Druid, HBase, Vertica.
- **Document:** Apache CouchDB, ArangoDB, BaseX, Clusterpoint, Couchbase, Cosmos DB, IBM Domino, MarkLogic, MongoDB, OrientDB, Qizx, RethinkDB
- **Key-value:** Aerospike, Apache Ignite, ArangoDB, Berkeley DB, Couchbase, Dynamo, FoundationDB, InfinityDB, MemcacheDB, MUMPS, Oracle NoSQL Database, OrientDB, Redis, Riak, SciDB, SDBM/Flat File dbm, ZooKeeper
- **Graph:** AllegroGraph, ArangoDB, InfiniteGraph, Apache Giraph, MarkLogic, Neo4J, OrientDB, Virtuoso



<http://hbase.apache.org>

Data model (wide column store): sorted dictionary mapping key values to rows, where a row is a map of column families to column values. Column values are sparse, hence they can be seen as a mapping from column name to values.

Unlike a relational database, the names and format of the columns can vary from row to row in the same table.

Query language: API's, basically a get(startkey, endkey) method to retrieve a range of rows.

An SQL-like language in [Apache Phoenix](#)



<https://hive.apache.org/>

Hive: data warehouse software to read, write, and manage large datasets residing in distributed storage using SQL.

Data model: relation with basic and complex value types. Complex types include: *structs, maps, and arrays*.

Query language: SQL dialect, including ROLLUP and CUBE and analytic functions. It also includes materialized views, partitioning, and columnar file formats.

It can access (read/write) HBase tables.

Hive is then another DW management system!
Conceptual and logical design remain the same!
Physical design is specific of Big Data platform!

Other Hadoop tools

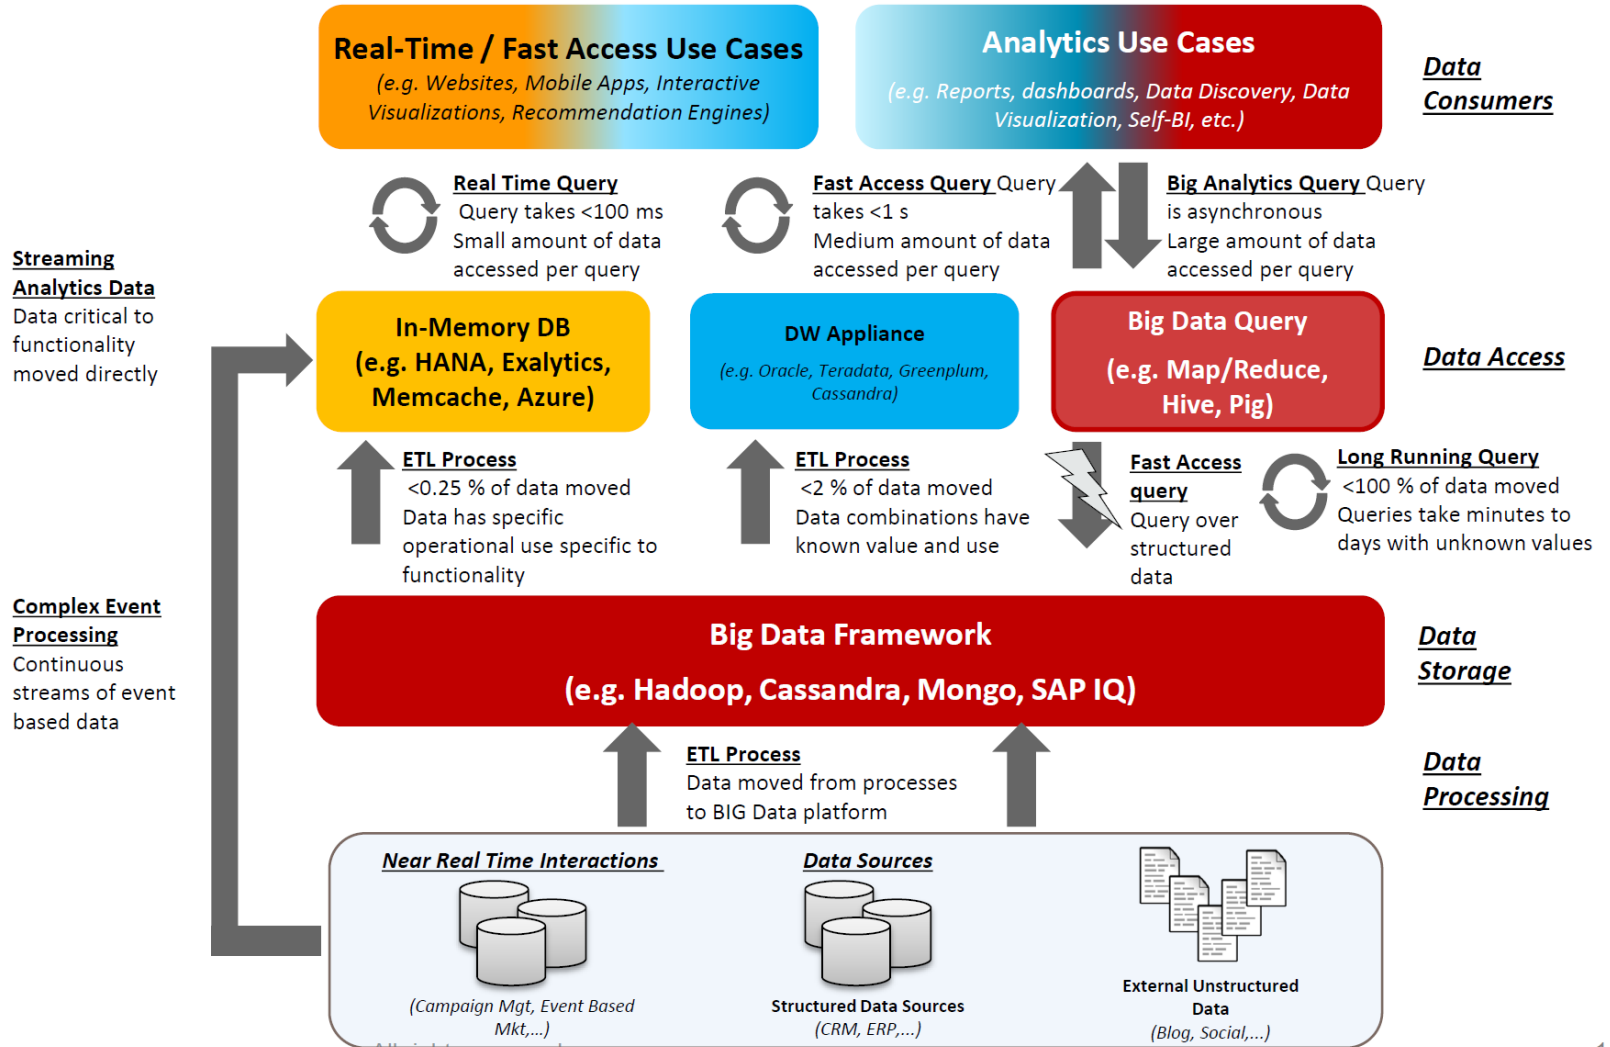
Sqoop: efficiently transferring bulk data between Apache Hadoop and structured datastores (RDBMS)



Other ways to query Hadoop with SQL:

- Oracle Big Data SQL
- SQL Server Big Data Clusters
- IBM BigSQL
- SparkSQL
- Hortonworks Stinger
- Apache Drill
- Cloudera Impala
- Facebook Presto
- ...

Enterprise Big Data Platform for DW



Summary of related courses

Distributed data analysis and mining (1st sem)

- Big data architectures and programming

Advanced databases (2nd sem)

- DBMS internals, query optimization

Visual analytics (2nd sem)

- Advanced reporting and storytelling

Technologies for web marketing (2nd sem)

- DSS for web analytics

Business process modeling (1st sem)

- Formal models for complex event processing

Legal issues in data science (2nd sem)

- Privacy, security and IPR in data management

Structure of exams

Written part (2 hours)


- See website for

- Examples of written text
- Dates and registration

- Grading ≥ 17 admission to the oral part

Mandatory teaching material

- [DW] A. Albano, S. Ruggieri. [Decision Support Databases Essentials](#), University of Pisa, 2 December 2020.
- [DB] A. Albano. [DB Essentials](#) and [solutions to exercises](#), University of Pisa, 1 December 2020. (English) from the book [Fondamenti di basi di dati](#) (in Italian, free download).
- Examples of [written exams with solutions](#) and [written exam](#).



Oral part

- Discussion of written part
- Open questions on all topics of the course
- Questions/small exercises using JRS and/or SQL Server

Time for filling student's questionnaires

<https://esami.unipi.it>



The screenshot shows the VALUTAMI website interface. The header features the university logo and the text "VALUTAMI VALUTazione della didattica ed iscrizione agli esAMI". The left sidebar menu includes items: Home, Exams, Subscribe, Subscriptions, Calendar by course, Calendar by session, Evaluation (highlighted with a red box and a red arrow), Syllabi, and Support. The main content area contains a light blue box with the following text: "Dal 25/01 e' disponibile il questionario relativo ai servizi e all'organizzazione della didattica. Tutti gli studenti sono inviati a partecipare esprimendo la propria opinione. Il questionario è anonimo, si compone di 13 domande e deve essere compilato una sola volta. Vai alla pagina".