

ORACLE[®]

Oracle Database 11g for Data Warehousing

Hermann Bär - Director Product Management, Data Warehousing



Oracle DW Strategy

- Best Database for BI/DW
 - 30 years of innovation
 - No other database can compare on the breadth and sophistication of Oracle's database features
- Within complete solutions
 - Complete database platform capabilities: ELT and Analytics
 - Complete BI and Performance Management solutions from Oracle
 - Broadest array of third-party technologies and solutions
- On the right hardware infrastructure

Complete Solutions


Performance Mgmt: Hyperion

BI Applications: Bus Obj, Cognos

BI Tool: Business Objects, Cognos

Middleware: IBM, BEA

Database: **ORACLE®**




Performance Mgmt: **ORACLE®**

BI Application: **ORACLE®**

BI Tool: **ORACLE®**

Middleware: **ORACLE®**

Database: **ORACLE®**



Benefits:

- Integrated stack
- Continued best-of-breed
- Top-to-bottom performance optimization

Oracle Data Warehouse Platform

ORACLE[®] 11^g
DATABASE

Analytics

Multi-Dimensional Calculations

Time Series

Statistics

Data Mining

Intelligent Query Processing

Materialized Views

Bitmap Indexes

Partition Elimination

Star Query Optimization

Scalable Data Management

Parallel Execution

Partitioning

RAC

Automatic Storage Mgmt

Compression

Data Integration

Bulk ETL

Real-Time ETL

Data Quality

Workload Management

Metadata Management

Security

Oracle10g for Data Warehousing

Continuous Innovation

Oracle 7.3

Oracle 8.0

Oracle8i

Oracle9i

Oracle10g

- Partitioned Tables
- Partitioned Indexes
- Parallel Indexes
- Parallel Indexes
- Parallel Bitmap Indexes
- Parallel Aggregate Functions
- Parallel Cost-Based Optimization
- Server Management
- Point-in-Time Recovery

- Hash and Join
- Resource Management
- Progress
- Adaptive
- Server-based
- Materialized Views
- Transportable Tablespace
- Direct Loading
- Functional Indexes
- Partitioned Tables
- Security

- List and Range
- Table Compression
- Bitmap Joins
- Self-Tuning
- New Analytic
- Grouping Sets
- External Tables
- MERGE
- Multi-Table Joins
- Proactive Query
- System Management

- Self-tuning SQL Optimization
- SQL Access Advisor
- Automatic Storage Manager
- Self-tuning Memory
- Change Data Capture
- SQL Models
- SQL Frequent Itemsets
- SQL Partition Outer Joins
- Statistical functions
- and much more ...

Oracle for Data Warehousing

Continuous Innovation

Oracle 7.3

Oracle 8.0

Oracle8i

- Hash Partitioning
- Resource Manager
- Program Partitioning
- Adaptive Server Partitioning
- Function-Based Partitioning
- Materialized Views
- Transparent Database Gateway
- Direct Path
- Functional Indexing
- Partitioned Security

Oracle9i

- List and Range Partitioning
- Table Compression
- Bitmap Indexing
- Self-Tuning SQL
- New Analytic Functions
- Grouping Functions
- External Tables
- MERGE Statement
- Multi-Table Joins
- Proactive System

Oracle10g

- Self-tuning SQL
- SQL Access Advisor
- Automatic Statistics
- Self-tuning SQL
- Change Data Capture
- SQL Model
- SQL Frequency
- SQL Partitioning
- Statistical Feedback
- and much more

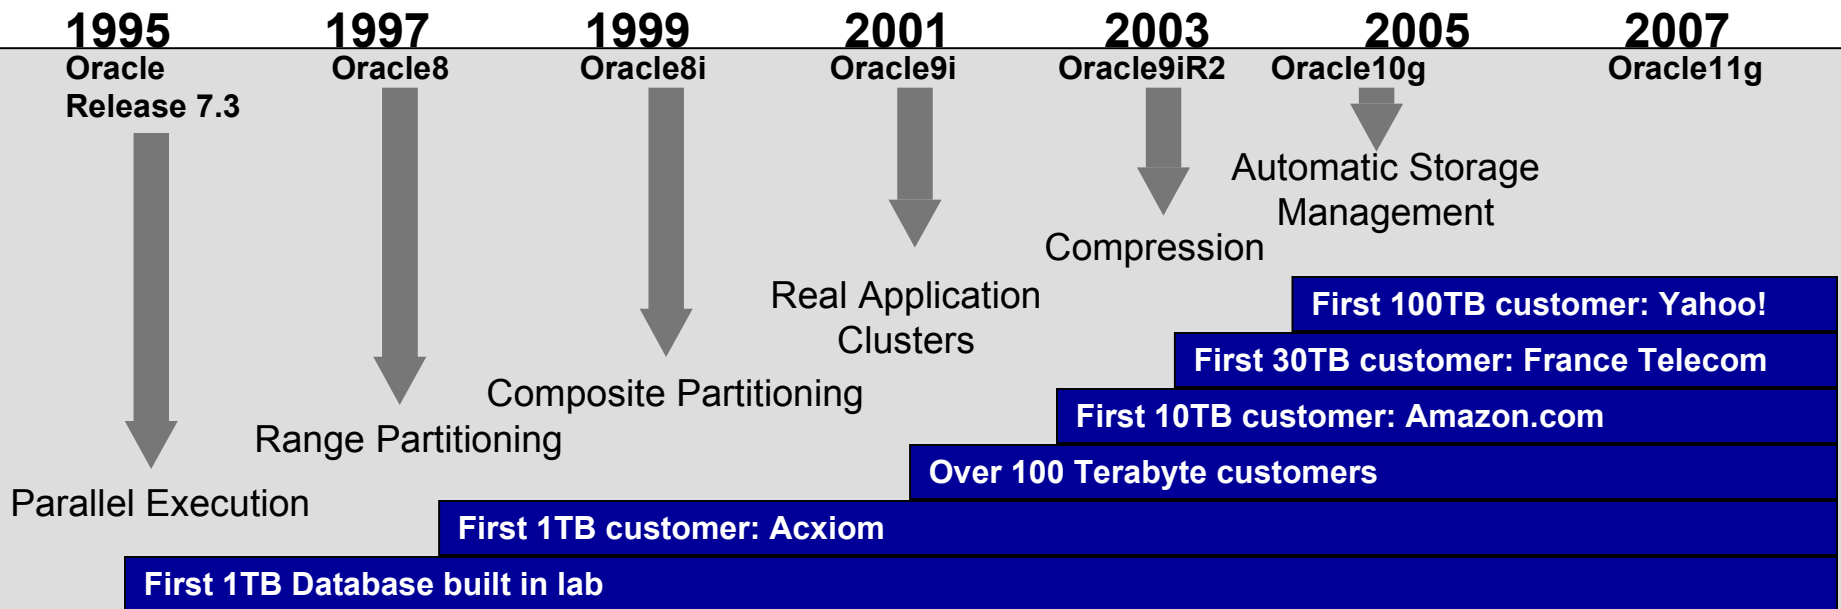
Oracle11g

- New composite partitionings
- Virtual column partitioning
- REF partitioning
- Cube-based Materialized Views
- SQL Pivot and Unpivot
- Query Result Cache
- SQL Plan Management
- General Linear Models
- Advanced Compression Option
- OWB included with DB

Continuous R+D Investment in VLDW


Continuous Customer Success in VLDW

- Over the past 12+ years, Oracle has steadily introduced major architectural advances for large database support
- Data warehouses have grown exponentially with these new technologies

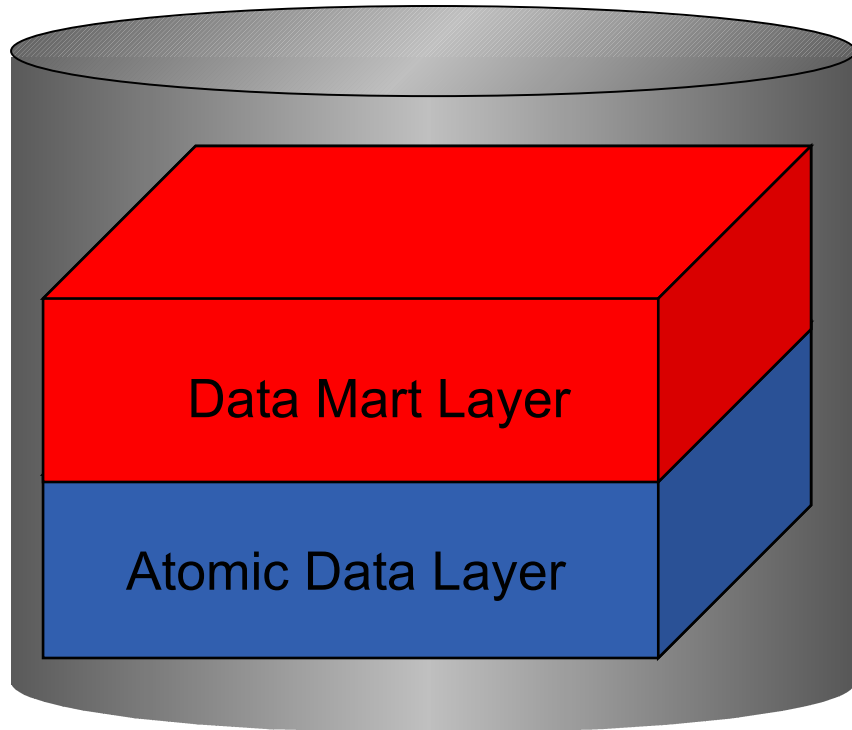


Long-standing dichotomy in the DW Market

- “Big brain” Workloads
 - Sophisticated database optimization techniques
 - Advanced Indexing
 - Dimensional query optimizations
 - Materialized views
 - Partition pruning
 - Algorithms and access paths determine performance
- Powerlifting workloads
 - Brute-force query execution
 - Large amounts of hardware
 - Query parallelism, hash partitioning
 - Hardware capabilities determine performance



Established Architectural Solutions



Application-specific performance structures
Summary data / materialized views
Dimensional view of data
Supports specific end-users, tools, and applications




Base data warehouse schema
Atomic-level data, 3nf design
Supports general end-user queries
Data feeds to all dependent systems



Half-Solutions are not the answer

- Enterprise DW solutions must provide both pieces
 - A solution that only provides one part will be limited to simple applications and unable to support growing
 - “If all you have is a hammer ...”



“Big Brain” Strategy



Oracle aggressively continues “Big Brain” strategy

- VLDB
 - Composite Range-Range
 - Composite List-Range
 - Composite List-List
 - Composite List-Hash
 - REF Partitioning
 - Virtual Column Partitioning
 - Compression enhancements
- Performance
 - Query Result Cache
- Manageability
 - Partition Advisor
 - Interval Partitioning
 - SQL Plan Management
 - Automatic SQL Tuning with Self-Learning Capabilities
 - Enhanced Optimizer Statistics Maintenance
 - Multi-Column Optimizer Statistics
 - ASM Fast Resync, Fast VLDB Startup and other enhancements
- SQL
 - SQL Pivot and Unpivot
 - Continuous Query Notification
- OLAP
 - Materialized view refresh and SQL rewrite
 - Continued database integration
 - Cube metadata in the Data Dictionary
 - Fine-grained data security on cubes
 - Simplified application development
 - Fully declarative cube calculations
 - Cost-Based Aggregation
 - Simpler calculation definitions
- Data Mining
 - Simplified development and deployment of models
 - Supermodels: data preparation combined with mining model
 - Additional packaged predictive analytics
 - Integration in database dictionary
 - New algorithms: “General Linear Models”
 - Encapsulates several widely used analytic methods
 - Multivariate linear regression; logistic regression
- ETL
 - OWB Repository installed with Database by default
 - Seibel connector
 - Graphical creation of views, materialized views

Oracle Partitioning:

Ten Years of Development

	Core functionality	Performance	Manageability
Oracle8	Range partitioning Global range indexes	“Static” partition pruning	Basic maintenance operations: add, drop, exchange
Oracle8i	Hash and composite range-hash partitioning	Partition-wise joins “Dynamic” pruning	Merge operation
Oracle9i	List partitioning		Global index maintenance
Oracle9i R2	Composite range-list partitioning	Fast partition split	
Oracle10g	Global hash indexes		Local Index maintenance
Oracle10g R2	1M partitions per table	“Multi-dimensional” pruning	Fast drop table
Oracle Database 11g	More composite choices REF Partitioning Virtual Column		Interval Partitioning Partition Advisor

Partitioning in Oracle Database 11g


Virtual Column-Based Partitioning

ORDERS

ORDER_ID	ORDER_DATE	CUSTOMER_ID...
9834	12-JAN-2007	65920
8300	14-FEB-2007	39654
3886	16-JAN-2007	4529
2566	19-JAN-2007	15327
3699	02-FEB-2007	18733

REGION	AS (SUBSTR(ORDER_ID, 6, 2))
US	US
EU	EU
EU	EU
US	US
US	US



- REGION requires **no storage**
- Partition by ORDER_DATE, REGION



Partitioning in Oracle Database 11g


Complete Composite Partitioning

- Range – range
- List – list
- List – hash
- List – range




RANGE-RANGE

Order Date by
Order Value



LIST-RANGE

Region by
Order Value



LIST-LIST

Region by
Customer Type

Advanced Compression

Compress Large Application Tables


- Transaction processing, data warehousing

Compress All Data Types

- Structured and unstructured data types

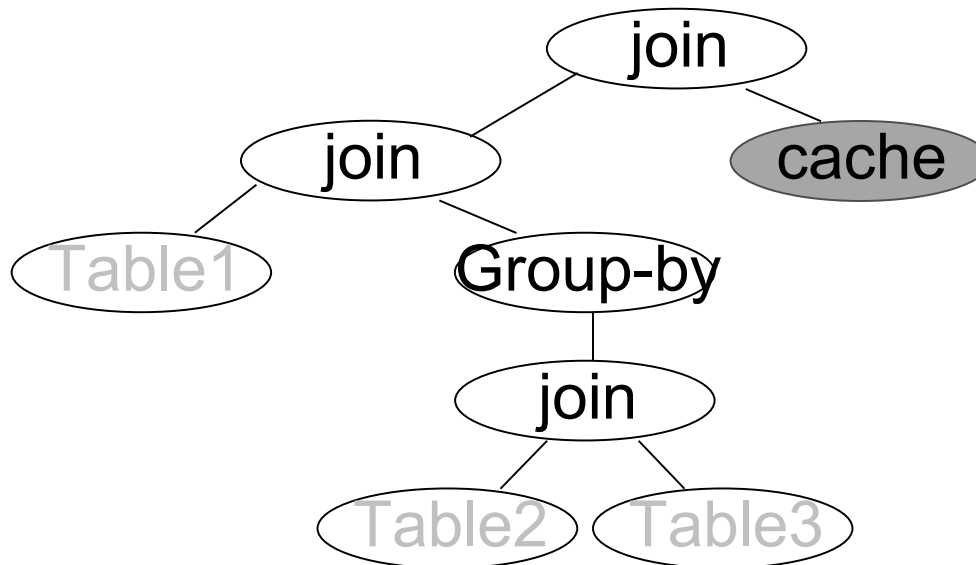
Typical Compression of 2-4X

- Cascade storage savings throughout data center



Database Result Cache


- Automatically Caches results of queries, query blocks, or pl/sql function calls
 - Cache is **shared** across statements and sessions on server
 - Significant speed up for read-only / read-mostly data
 - Full consistency and proper semantics
 - Cache refreshed when any underlying table updated



Q2: Uses result transparently

SQL Query Result Cache

- Retail customer data (~50 GB)
- Concurrent users submitting queries randomly
 - executive dashboard application with 12 heavy analytical queries
- Cache results *only* at in-line view level
 - 12 queries run in random, different order with 4 queries benefiting from the cache



Transparent “Big Brain” Features

- Materialized Views


- Transparent rewrites of expensive queries
 - Including rewrites on remote objects
- Incremental automatic refresh

ORACLE
DATABASE 11g



- Bitmap Indexes

- Optimal storage
- Ideal for star or star look-a-like schemas




- SQL Access Advisor – based on workload

- Materialized view advice
- Index advice
- Partition advice

ORACLE
DATABASE 11g



Integrated Analytics



- Bring the analytics to the data
- Leverage core database infrastructure

Native Support for Pivot and Unpivot

SALESREP	Q1	Q2	Q3	Q4	SALESREP	QU	REVENUE
100	230	240	260	300	100	Q1	230
101	200	220	250	260	100	Q2	240
102	260	280	265	310	100	Q3	260
					100	Q4	300
					101	Q1	200
					101	Q2	220
					101	Q3	250
					101	Q4	260
					102	Q1	260
					102	Q2	280
					102	Q3	265
					102	Q4	310

Native Support for Pivot and Unpivot

QUARTERLY_SALES

SALESREP	Q1	Q2	Q3	Q4
100	230	240	260	300
101	200	220	250	260
102	260	280	265	310

SALESREP	QU	REVENUE
100	Q1	230
100	Q2	240
100	Q3	260
100	Q4	300
101	Q1	200
101	Q2	220
101	Q3	250
101	Q4	260
102	Q1	260
102	Q2	280
102	Q3	265
102	Q4	310

```

select * from quarterly_sales
unpivot include nulls
(revenue for quarter in (q1,q2,q3,q4))
order by salesrep, quarter ;

```

Native Support for Pivot and Unpivot

SALES_BY_QUARTER

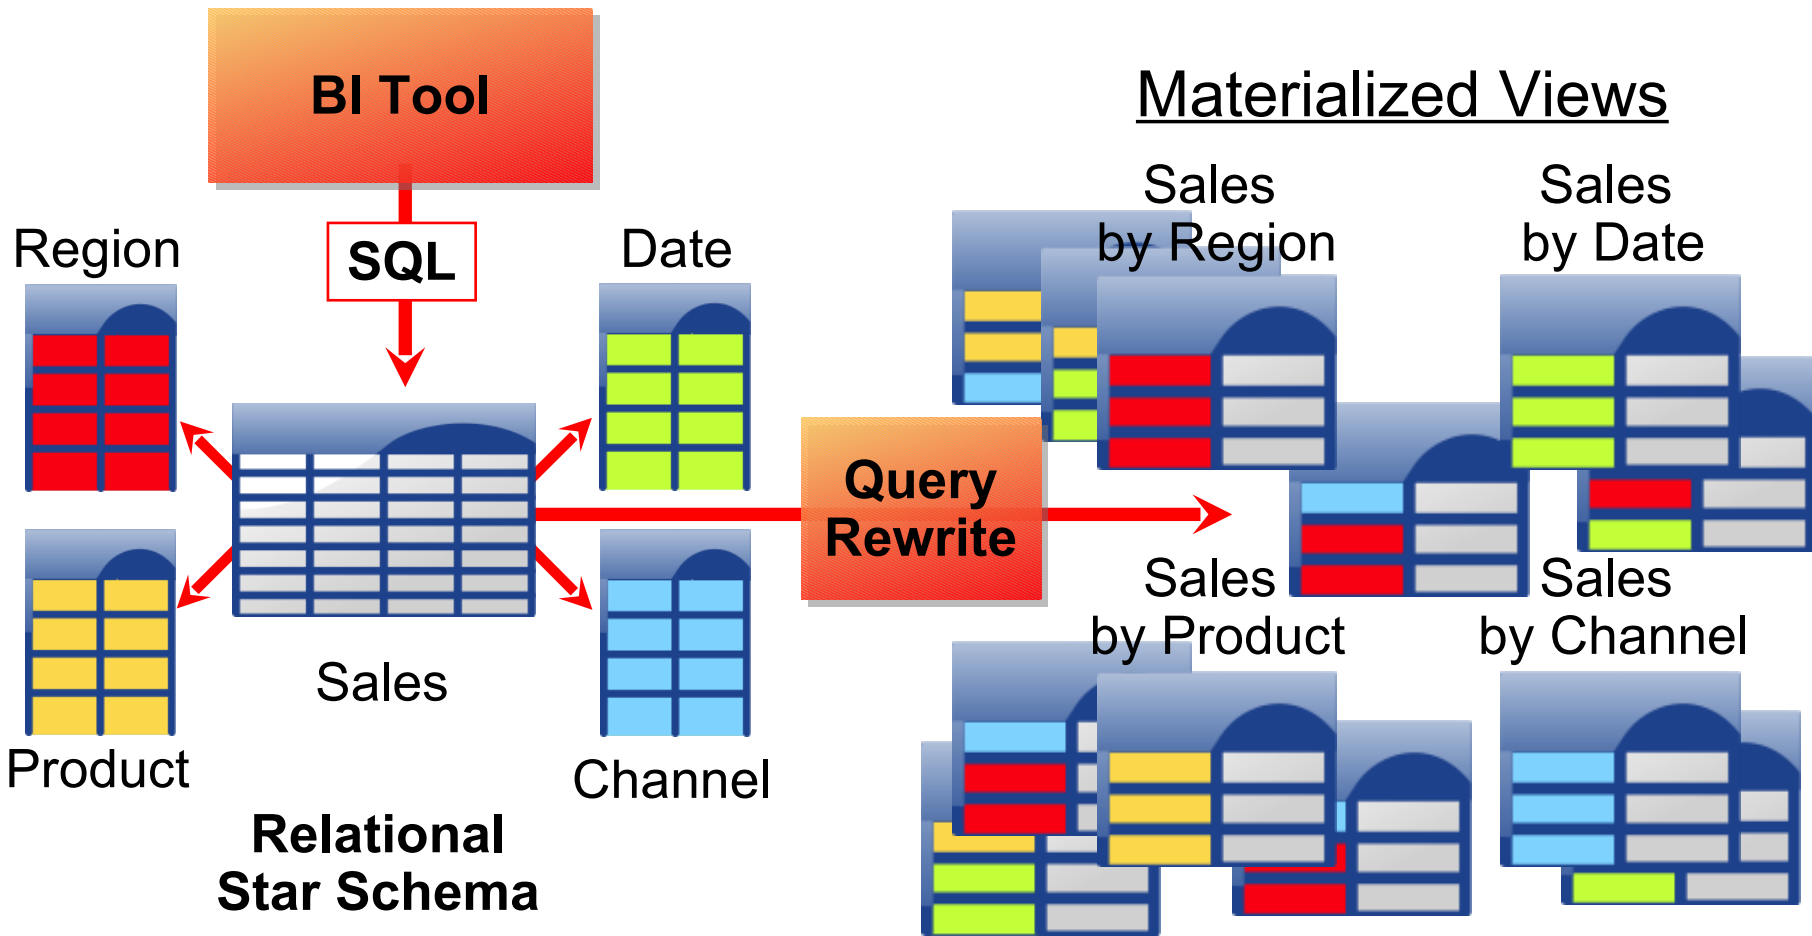
SALESREP	'Q1'	'Q2'	'Q3'	'Q4'
100	230	240	260	300
101	200	220	250	260
102	260	280	265	310

SALESREP	QU	REVENUE
100	Q1	230
100	Q2	240
100	Q3	160
100	Q4	90
100	Q3	100
100	Q4	140
100	Q4	70
101	Q1	200
101	Q2	220
101	Q3	250
101	Q4	260
102	Q1	260

```
select * from sales_by_quarter
pivot (sum(revenue)
for quarter in ('Q1', 'Q2', 'Q3', 'Q4'))
order by salesrep ;
```

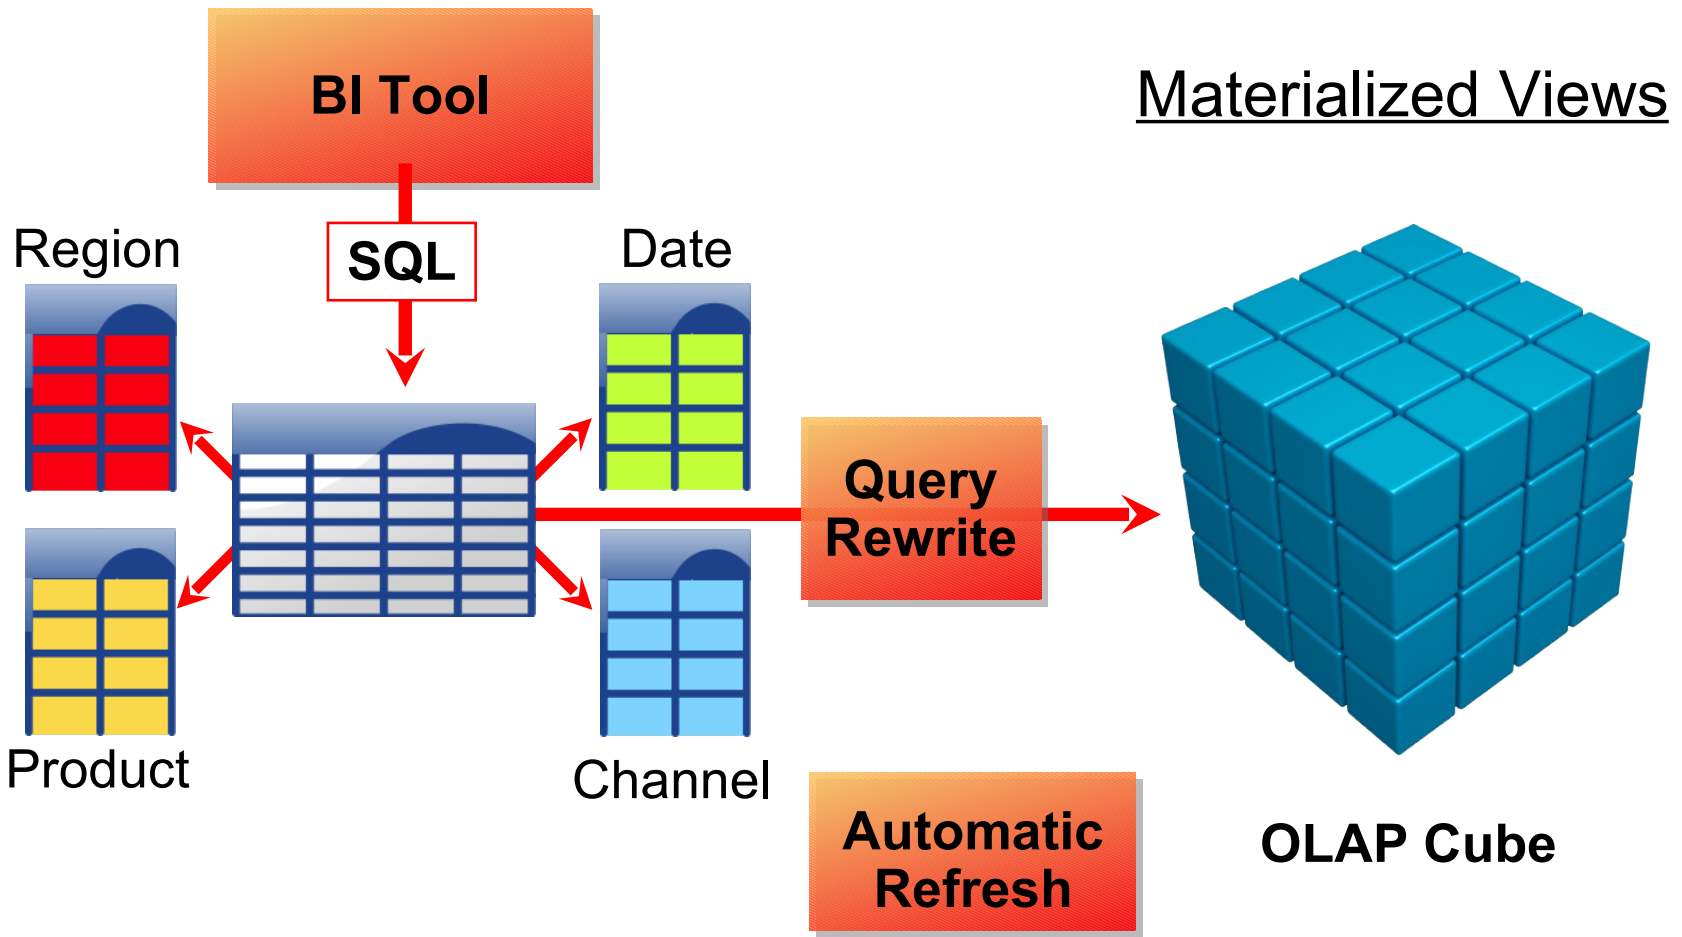
Business Intelligence Analysis

Typical Architecture Today



New in Oracle Database 11g

Cube-Organized Materialized Views



Oracle Warehouse Builder Packaging

Enterprise ETL Option

- Performance
- Productivity
- Reusability
- Metadata Management

Data Quality Option

- Data Profiling
- Anomaly Detection
- Business Rules
- Audit


ERP/CRM Connectors

- Oracle EBS
- PeopleSoft
- Siebel
- SAP

Core Features

(No extra cost with database SE/SE1/EE)

Self Managing Database



“Powerlifting” Strategy



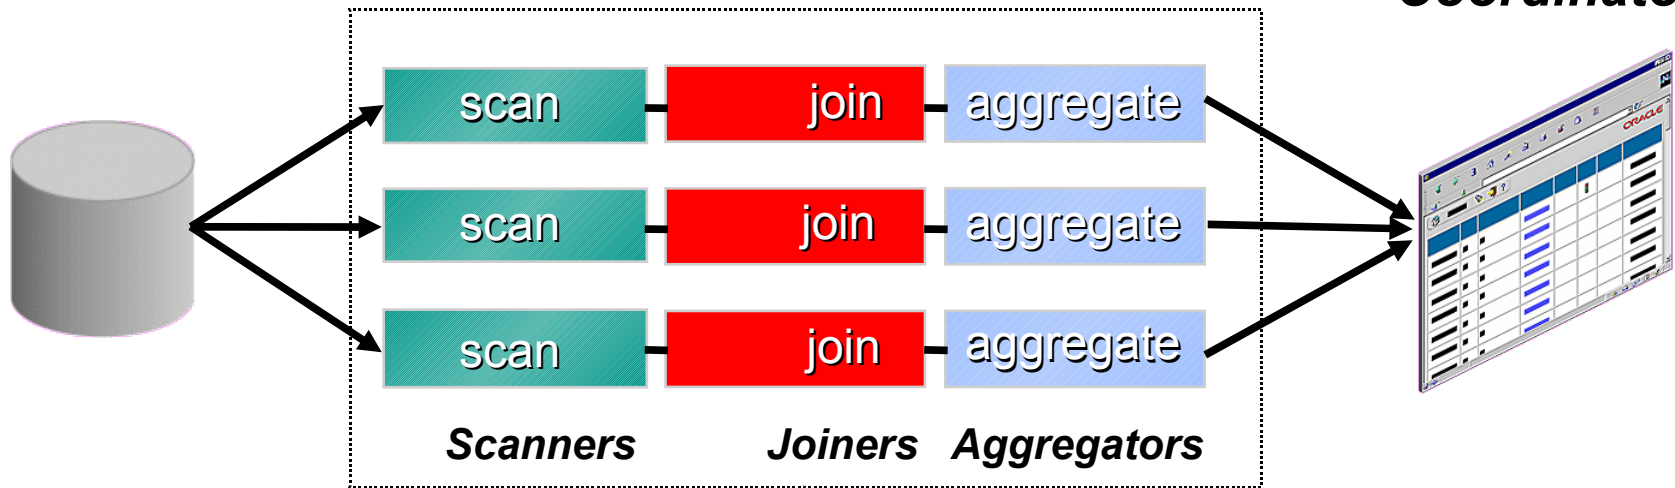
Parallel Execution

```
select c.cust_last_name  
, sum(s.amount_sold)  
from customers c, sales s  
where c.cust_id = s.cust_id  
group by c.cust_last_name ;
```

Data on Disk

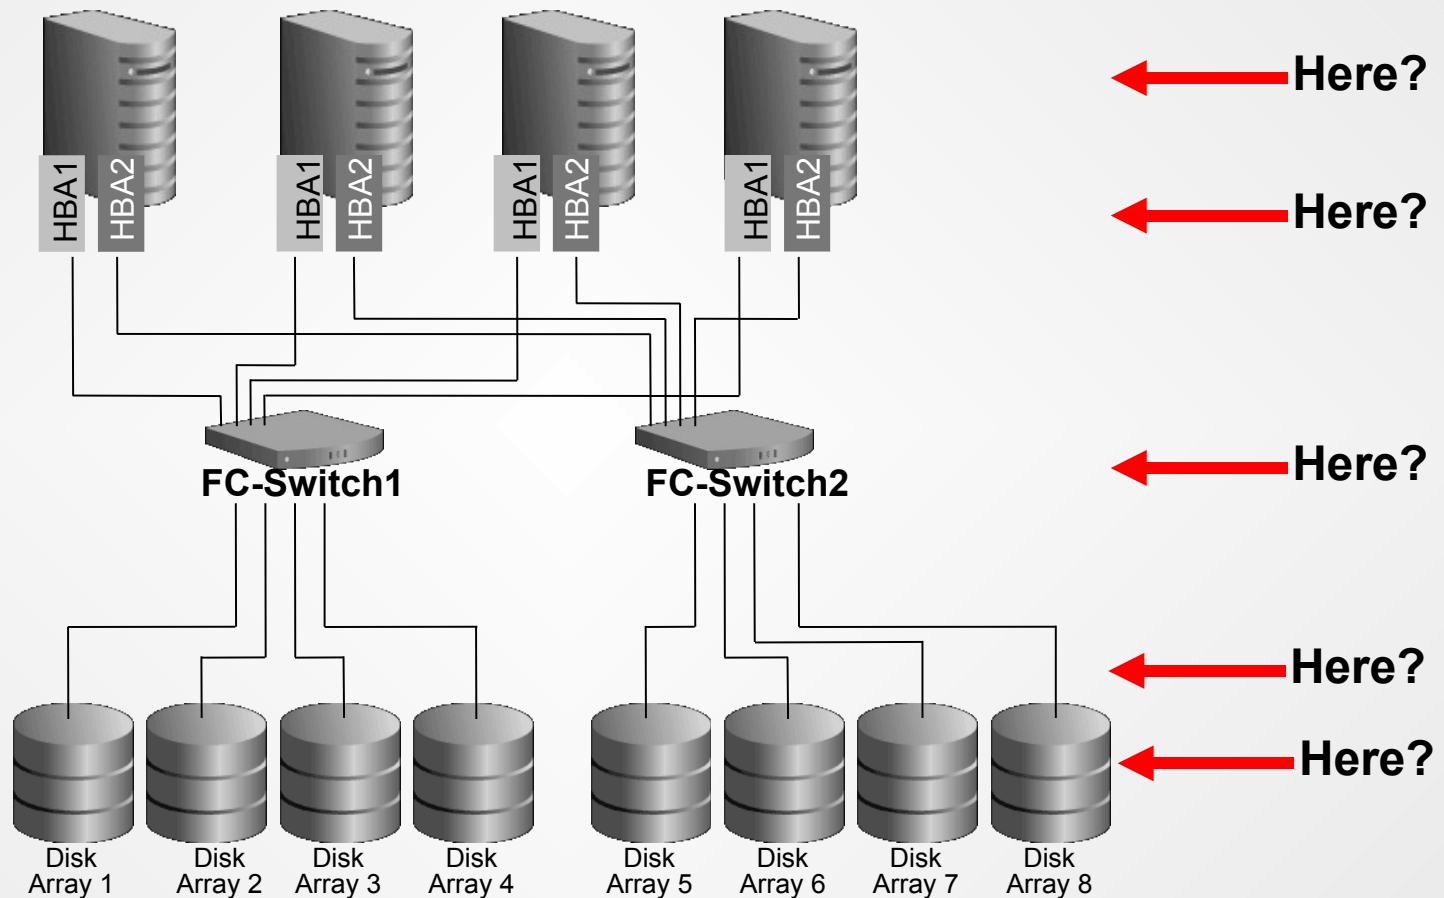
Parallel Servers

Coordinator



Typical Customer DW Platform

- Where is the performance bottleneck?



Only Balanced Configurations Drive Optimized Performance

An Unbalanced Configuration

100%
Possible
Efficiency

Database CPUs Memory Actuators LUNs Disks Raid

< 50%
Achieved
Efficiency

A Balanced Configuration

100%
Possible
Efficiency

Database CPUs Memory Actuators LUNs Disks Raid

100%
Achieved
Efficiency

Full Range of Options

Custom Solutions




Flexibility for the most demanding data warehouse

Reference Configurations



Documented best-practice configurations for data warehousing

Optimized Warehouse



Scalable systems pre-installed and pre-configured:
create-table ready

Pre-configured, Pre-installed, Validated

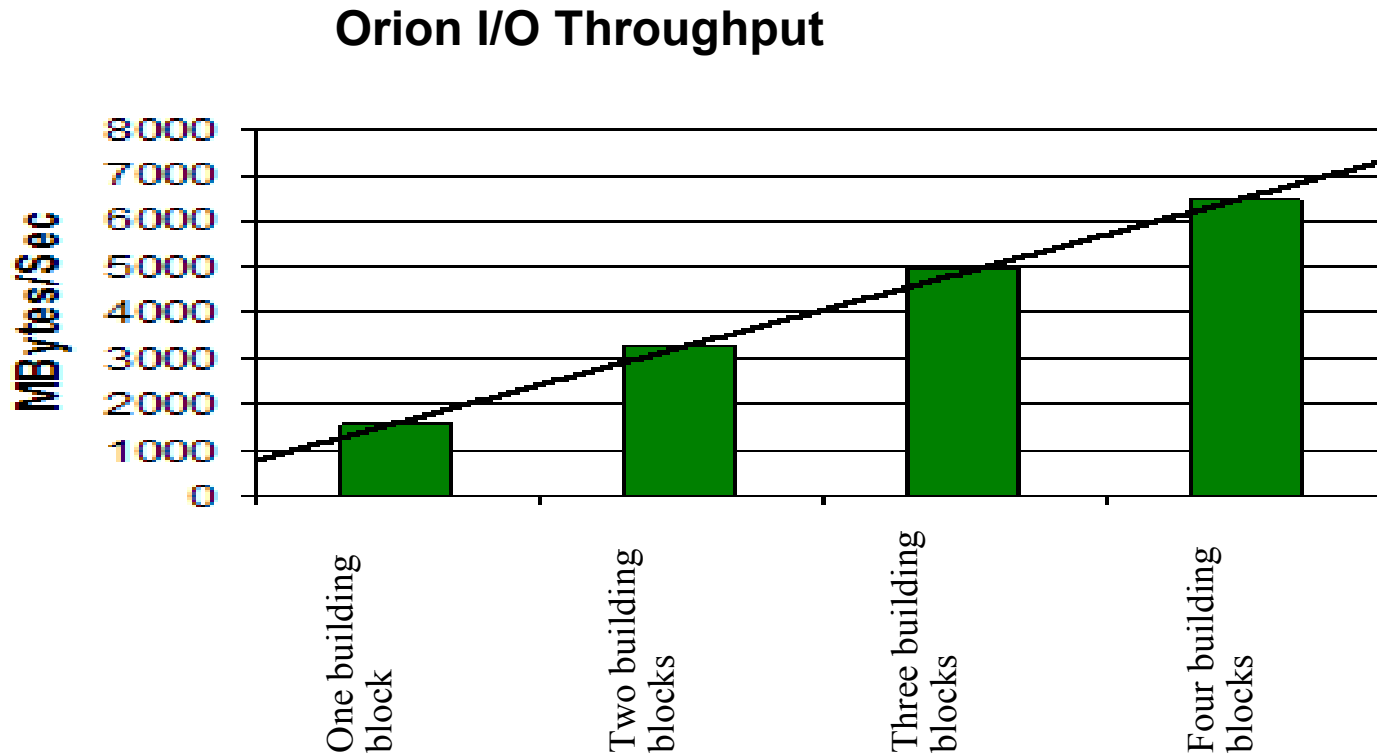
Complete Flexibility

Building Block Scale-Out

- Add power in a balanced fashion



Provide linear hardware scalability



Reduce Risk & Accelerate Deployment



Custom Configuration

Months.....

Testing and Validation

Installation and Configuration

Acquisition of Components

Pre-Implementation System Sizing

Weeks.....

Reference Configuration



Testing and Validation

Installation and Configuration

Acquisition of Components

Oracle Optimized Warehouse



Faster Deployment
Lower Risk
Increased Flexibility

Oracle Optimized Warehouse

< 1 Week to implement

Oracle Optimized Warehouse Availability



Scalability

Limited


High

Very High

Oracle DW Strategy

- Best Database for BI/DW
 - 30 years of innovation
 - No other database can compare on the breadth and sophistication of Oracle's database features
- Within complete solutions
 - Complete database platform capabilities: ELT and Analytics
 - Complete BI and Performance Management solutions from Oracle
 - Broadest array of third-party technologies and solutions
- On the right hardware infrastructure

#1 for Data Warehousing



**Worldwide Data Warehouse Management
Market Share, 2006**

Q & A



ORACLE®