



Emerging Innovations In Analytical Databases

TDWI India Chapter Meeting, July 23, 2011, Hyderabad

Vivek Bhatnagar

Agenda

- ❑ Today's Biggest Challenge in BI - SPEED
- ❑ Common Approaches Used Till Date for Performance
- ❑ New Breakthrough: On-Chip Vectorised Computing/
Columnar Design
- ❑ VectorWise Analytical Database

Today's Challenges in BI



Make informed decisions faster

Analysis in seconds not minutes, minutes not hours



Data explosion

Collecting more information, less resources
>44x growth in next 10 years



Existing Tools:

Too Slow. Too complex. Too expensive.

Analytical databases designed in 80s & 90s do not take advantage of today's modern hardware

Biggest BI Challenge - SPEED

What Problem will eventually drive you to replace your current data warehouse platform?

1. Poor Query Response **45%**

2. Can't Support Advanced Analytics **40%**

Source: TDWI Q4 2009 Best Practices Report

“Gartner clients increasingly report performance constrained data warehouses during inquiries. Based on these inquiries, we estimate that nearly 70% of data warehouses experience performance-constrained issues of various types.”

Gartner

Source: Gartner Magic Quadrant for Data Warehouse Database Management Systems, Jan 2010

Biggest BI Challenge - SPEED

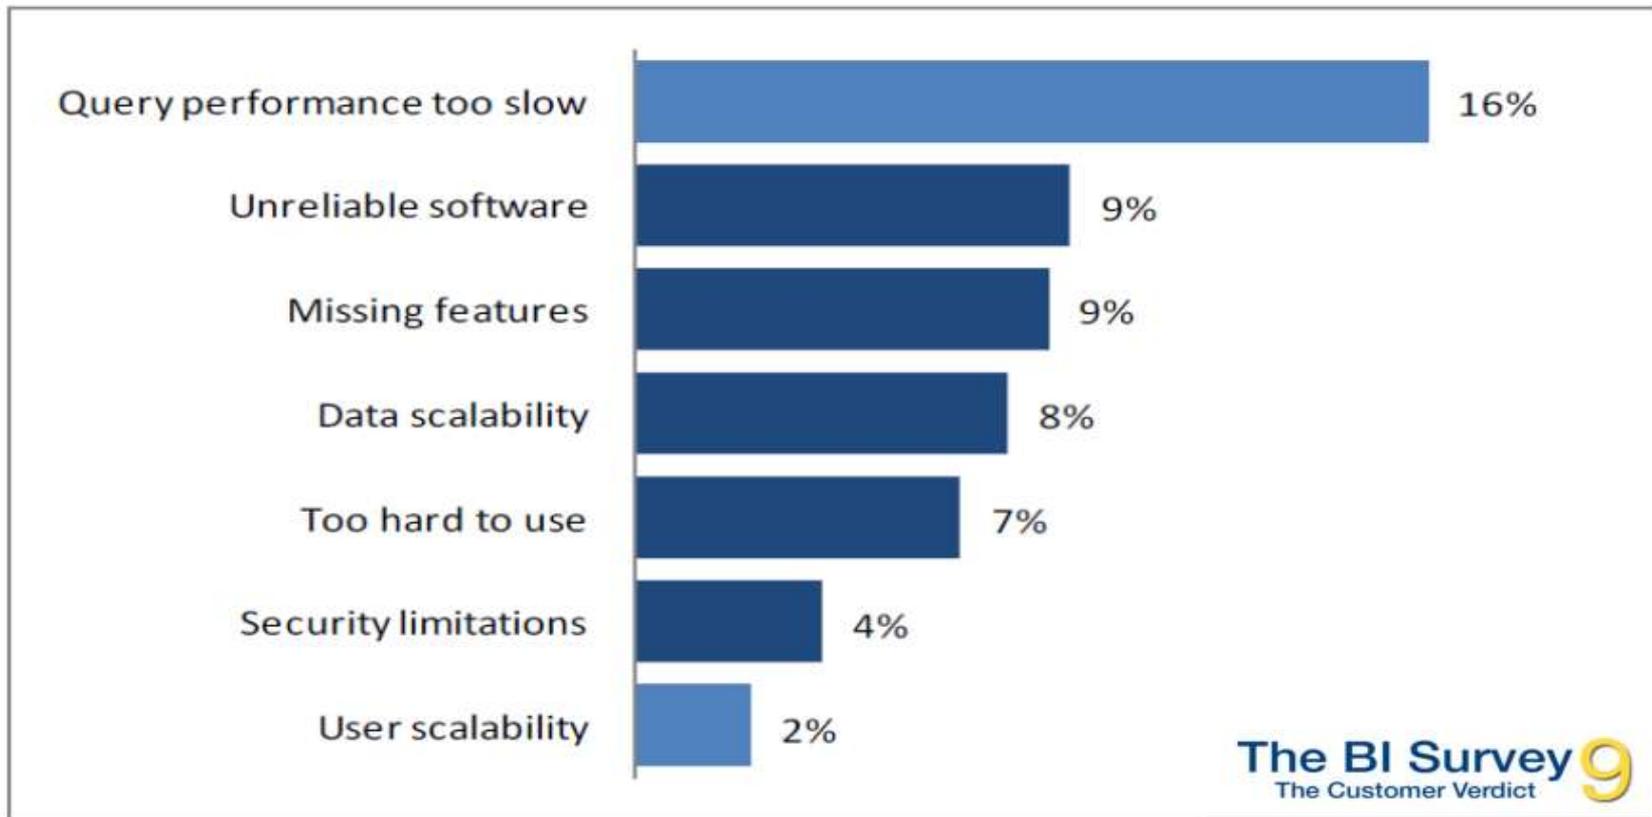


Figure 150: Performance problems (in lighter blue) compared to other product-related problems

Source: 2010 BI Survey 9 – World's largest independent BI Survey 3093 respondents

Cubes and Speed

As cubes grow in size they take longer to load and build

- Processing time might exceed batch window
- Difficulty managing large cubes
- Time required to add new dimensions

On average, how long does it take to add a new source of data to your data warehouse?



Average time to add a new data source	
2008	7.0 weeks
2009	8.4 weeks
2010	7.4 weeks

Source: 2010 TDWI Benchmarks

Relational Databases and Speed

Limitations in SQL technology

- Adhoc queries too slow
- Indexing/Aggregations cost time & money
- 25% average BI/DW team time used up for maintenance/change management

What percentage of your BI/DW team is allocated to these tasks?

	2008	2009	2010
Development/testing	52%	50%	52%
Maintenance/change management	25%	26%	25%
Support/training	16%	15%	14%
Other	9%	10%	11%

Source: 2010 TDWI BI Benchmark Report

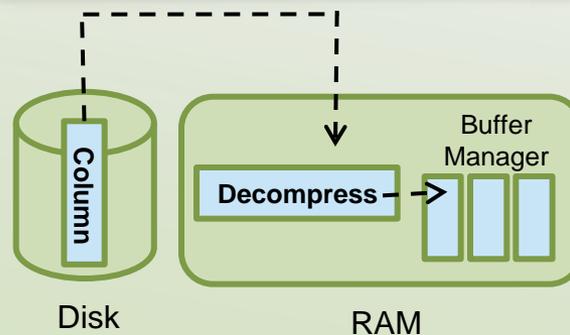
Common Approaches Used Till Date for Achieving Database Performance

Optimizations for parallel processing and minimal data retrieval

Parallel Processing



Column store with compression



Proprietary hardware



Data Warehouse Appliances

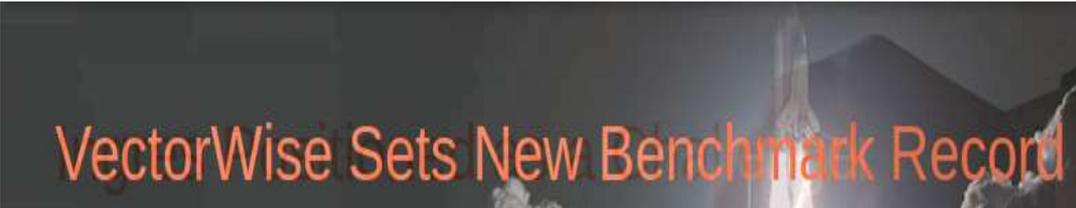
Acceptable performance has been achieved by using more hardware or by intelligently lowering the volume of data to be processed

However, none of these approaches leverages the performance features of today's CPUs i.e. taking the most out of each modern commodity CPU

New Breakthrough Analytical Database

Relational database for BI and data analysis

- Runs blazing fast/interactive data analysis
- Exploits performance potential in today's CPUs
- Delivers in-memory performance without being memory constraint



VectorWise Sets New Benchmark Record

“Game-changing technology.”
Don Feinberg, Gartner Group

“This is definitely a breakthrough. It delivers faster results at lower costs.”

Noel Yuhanna, Forrester Research

“This inevitability puts VectorWise 4 years ahead of the competition in terms of performance – and it will remain 4 years ahead until some competitor finds a way to catch up at a software level. This is unprecedented.”

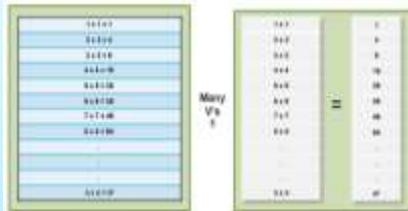
Robin Bloor, The Virtual Circle

On Chip Vectorised Computing/Columnar Database

Breakthrough technology

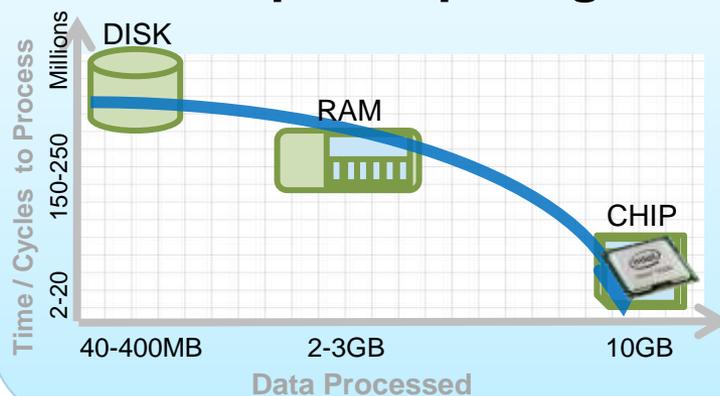
Innovations on industry-proven techniques

Vector Processing



Updateable Column Store

On Chip Computing



Automatic Compression



Automatic Storage Indexes

Minimize IO



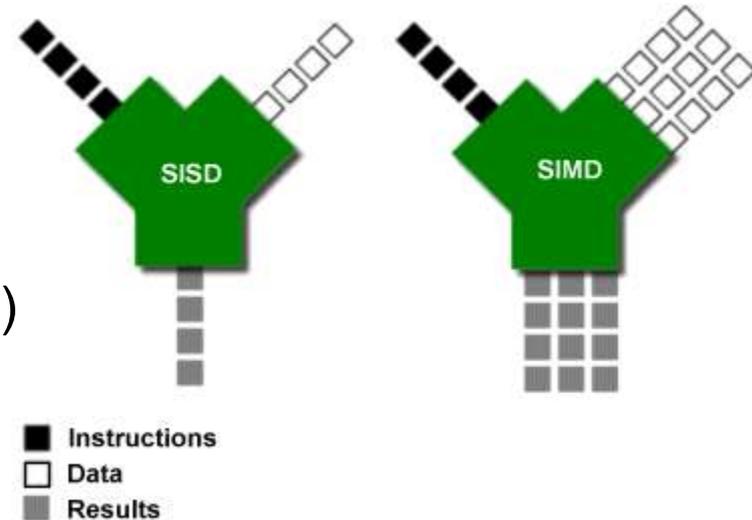
Parallel Processing

VectorWise Technology

- **Vector processing**
 - Exploits super-scalar features using SIMD capabilities of today's CPUs
- **Optimizes memory hierarchy**
 - Maximizes use of CPU cache
 - Fewer requests to RAM and disk
- **Data Compression/De-Compression**
 - Optimized compression enabling very fast de-compression for overall performance enhancement
 - Vectorized de-compression
 - Automatic compression through ultra-efficient algorithms
- **Automatic Indexing**
 - System generated Storage Indexes
 - Easy identification of candidate data blocks for queries
- **Integration**
 - Standard SQL and interfaces
 - Common BI/Data Integration tools

Modern CPU Instruction Capabilities

- **SIMD**
 - Traditional CPU processing: Single Instruction, Single Data (SISD)
 - Modern CPU processing capabilities: Single Instruction, Multiple Data (SIMD)
- **Out-of-order execution**
- **Chip multi-threading**
- **Large L2/L3 caches**
- **Streaming SIMD Extensions for efficient SIMD processing**
- **Hardware accelerated String Processing**



Vector Processing

Traditional Scalar Processing

$1 \times 1 = 1$
$2 \times 2 = 4$
$3 \times 3 = 9$
$4 \times 4 = 16$
$5 \times 5 = 25$
$6 \times 6 = 36$
$7 \times 7 = 49$
$8 \times 8 = 64$
.
.
.
$n \times n = n^2$

One operation performed on one element at a time

Large overheads

Vector Processing

1×1	$=$	1
2×2		4
3×3		9
4×4		16
5×5		25
6×6		36
7×7		49
8×8		64
.		.
.		.
.		.
$n \times n$		n^2

Many V's
1

One operation performed on a set of data at a time

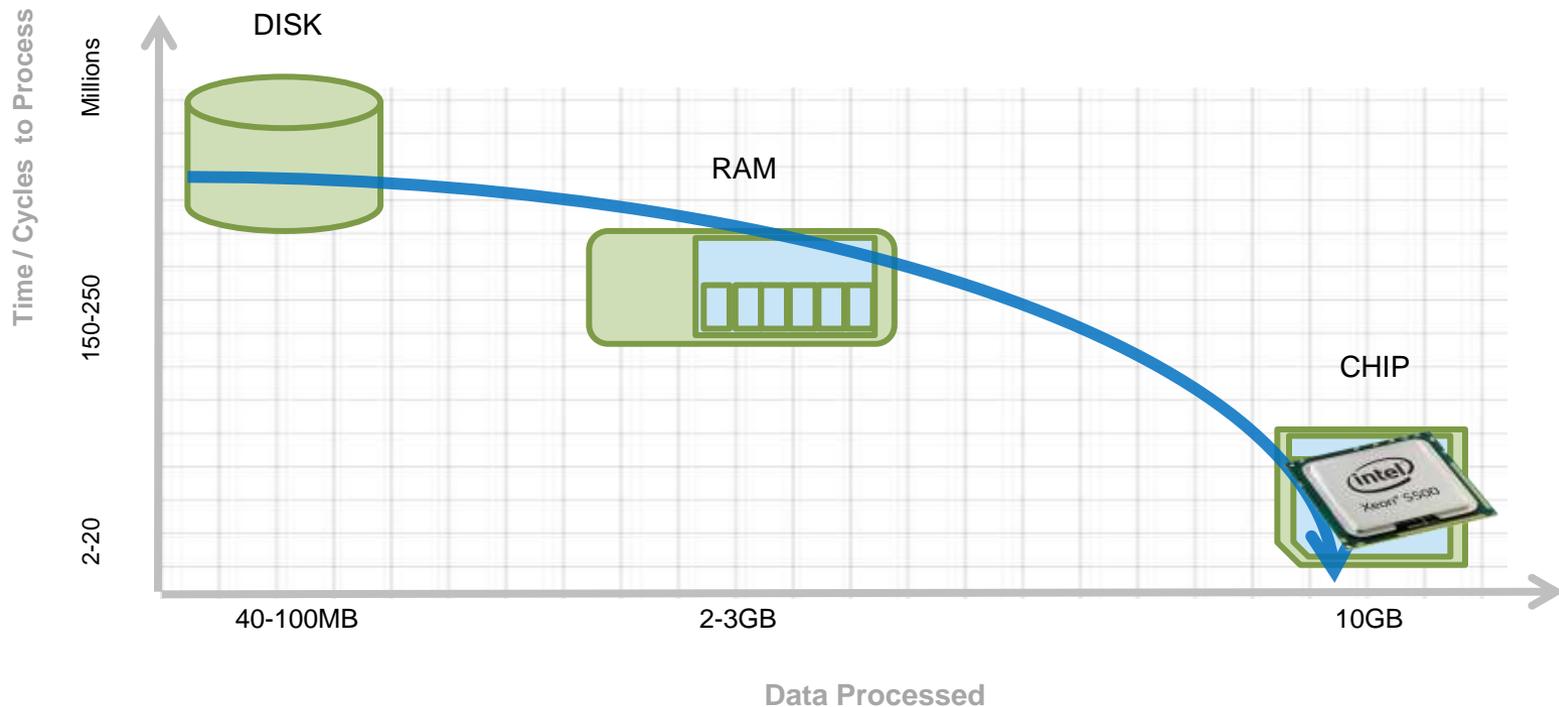
No overheads

Process even 1.5GB per second

Processing in Chip Cache

GB/s	Measure of Throughput
Cycles	Amount of CPU time required to process data

Using CPU cache is far more faster & efficient



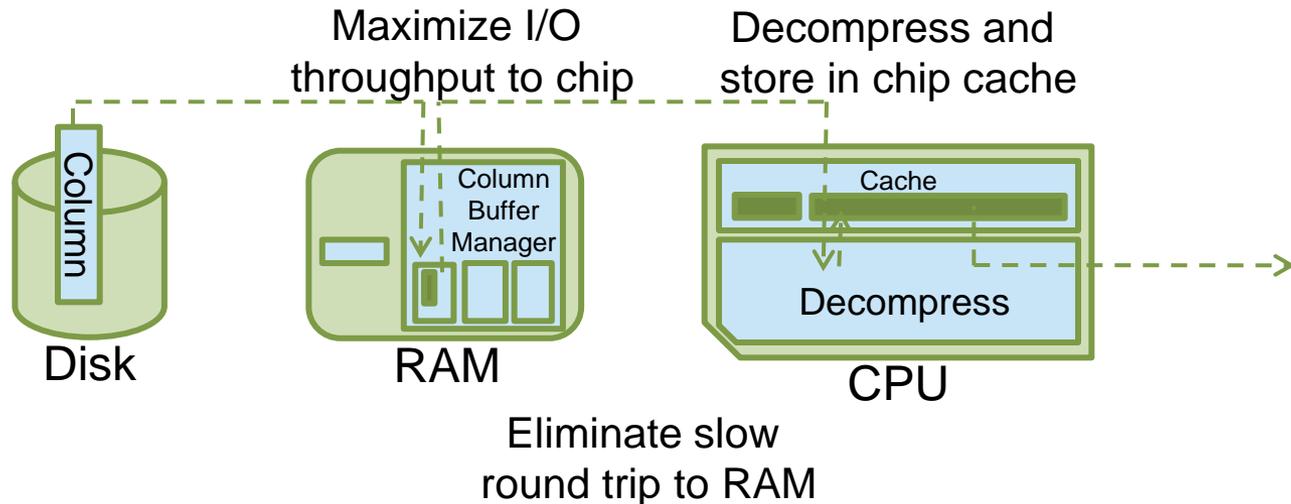
Updateable Column Store

- Only access relevant data
- Enable incremental updates efficiently
 - Traditionally a weakness for column-based stores

Cust_Num	Cust_surn_ame	Cust_first_na me	Cust_mid_na me	Cust_DOB	Cust_Sex	Cust_Add_1	Cust_Addr_2	Cust_City	Cust_State
46328927956	Jones	Steven	Sean	17-JAN-1971	M	333 StKilda Rd		Melbourne	Vic
98679975745	Smith	Leonard	Patrick	04-APR-1964	M	Unit 12, 147 Trafalgar Sqr		Birmingham	London
52634346735	Rogers	Cindy	Carmine	11-MAR-1980	F	Belmont Rail Service	421 Station St	Belmont	CA
346737347347	Andrews	Jenny		14-SEP-1977	F	Apt1, 117 West 42 nd St		New York	NY
88673477347	Cooper	Sheldon	Michael	30-JUN-1980	M	Ingres Corporation	Level 2, 426 Argello St	Redwood City	CA
34673447568	Kollwitz	Rolf		22-DEC-1975	M	IBM Headquarters	123 Mount View Crs	Atlantic City	PN
99554443044	Wong	Penny	Lee	13-NOV-1981	F	Ming On Tower 1	1777 Moa Tzu Tung Rd	Ming Now Province	Shanghi

Optimized Compression & Fast De-Compression

- **Column-based compression with multiple algorithms**
 - Automatically determined by Ingres VectorWise
- **Vectorised decompression**
 - Only for data processing in CPU cache



Storage Index

- **Always automatically created**
- **Automatically maintained**
- **Stores min/max value per data block**
- **Enables database to efficiently identify candidate data blocks**

VectorWise Features

Performance

- 10x-75x faster for BI, analytics & reporting
- In-memory performance without memory restraints
- Near real-time updatable database
- Delivers results in seconds not minutes
minutes not hours

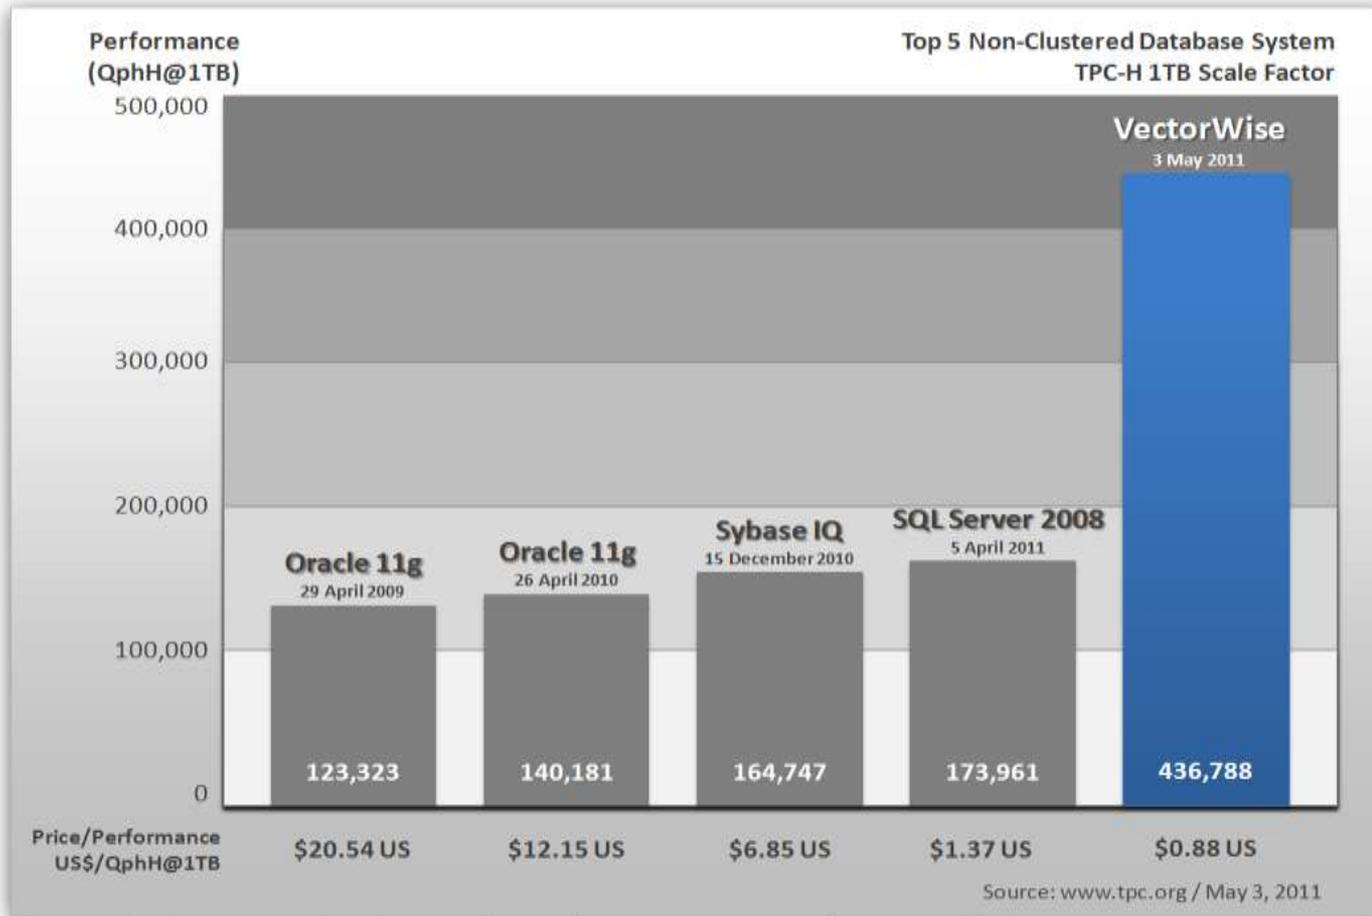
Usage & Integration

- Uses ANSI standard queries & SQL statements
- Eliminate/reduce Cubes, aggregate tables, roll ups, indexes....
- Self indexing & self tuning database
- Deliver BI projects faster with lower cost & risk

TCO

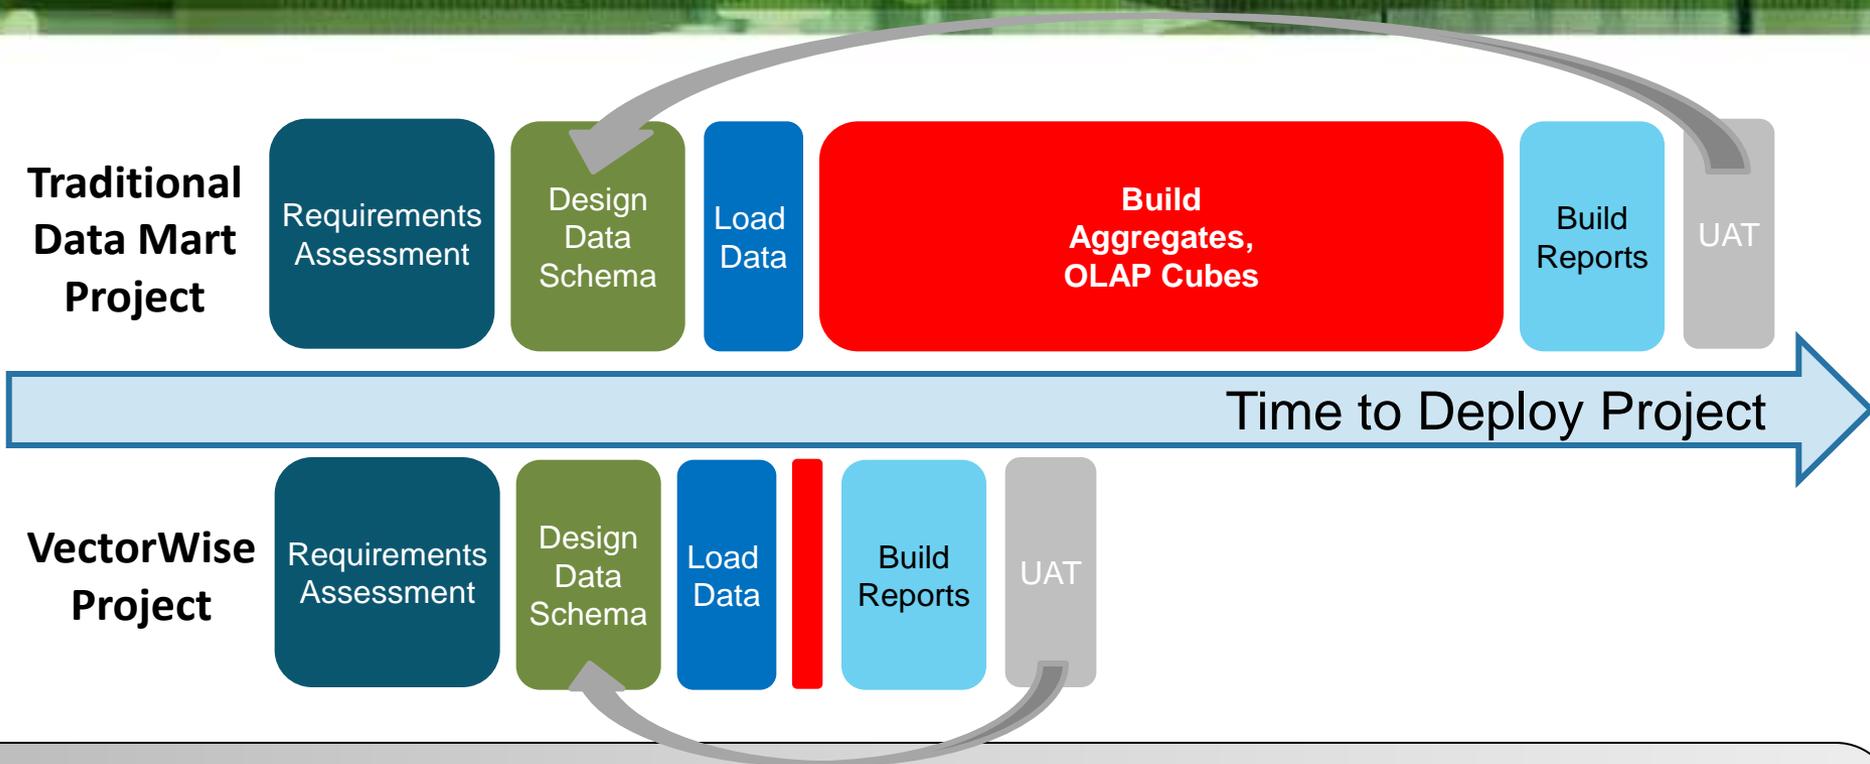
- Maximize utilization of CPUs in low cost commodity hardware
- Handle tens of terabytes scale data with a single server
- Requires commodity hardware
- Does not require MPP

TPC-H Performance Benchmark



TPC, TPC Benchmark, TPC-H, QppH, QthH and QphH are trademarks of the Transaction Processing Performance Council (TPC)

BI Tuning & Complexity



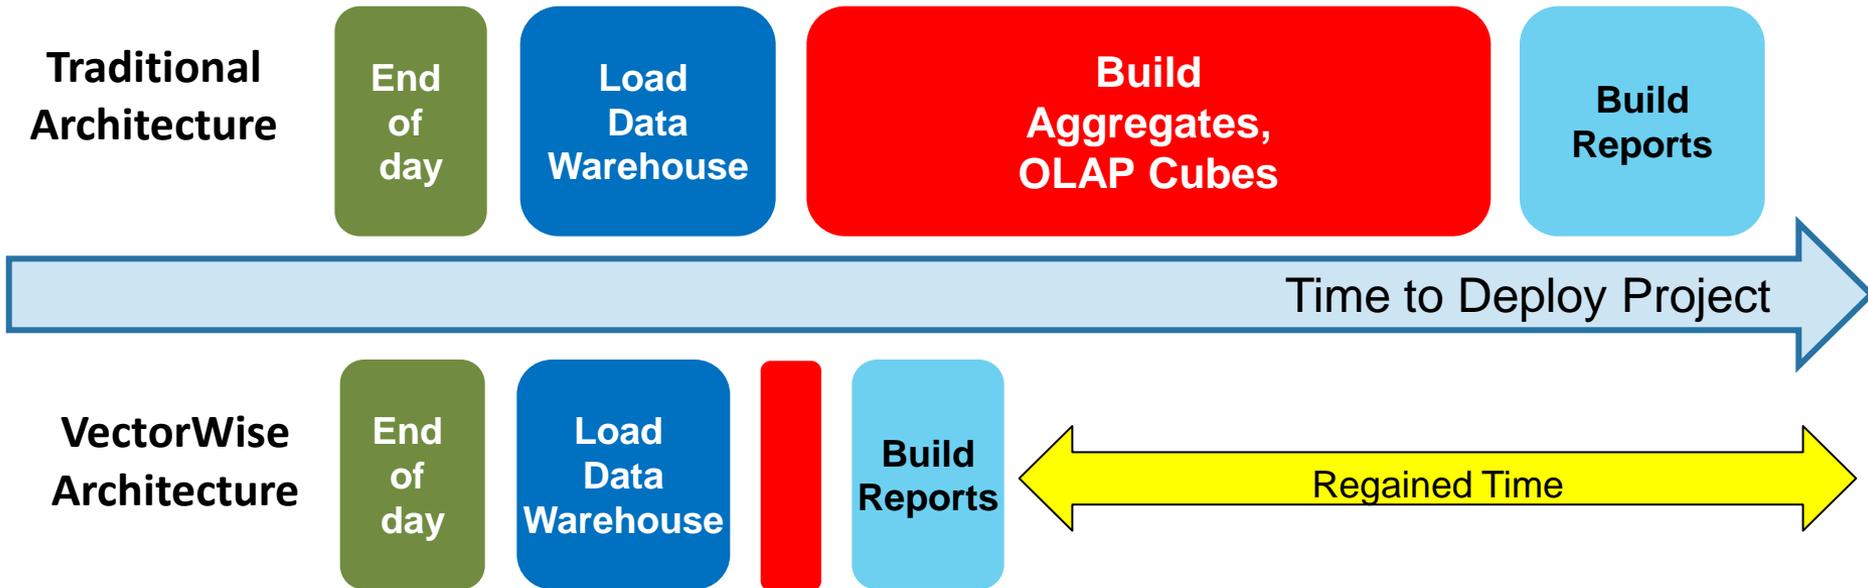
2010 TDWI BI Benchmark Report

Average time to build a complex report or dashboard
(20 dimensions, 12 measures, and 6 user access roles)

2008	6.7 weeks
2009	6.3 weeks
2010	6.6 weeks

BI Tuning & Complexity

Fast Processing Everyday!!!



TPC-H Price/Performance Benchmark

Price/Performance Benchmarks TPC-H 100GB Scale Factor

Non-Clustered Results - Price / Composite queries per hour (QphH)

Previous Benchmark 0.51 US\$/QphH

VectorWise 0.38 US\$/QphH

25.5% better price/performance than previous-benchmark record holder

Source: www.tpc.org / February 15, 2011

TCO

- Lower infrastructure cost
- Lower BI tuning costs