



Ingres VectorWise

Peter Boncz – co Founder

INGRES

vectorwise

INGRES

Analysts Twitter



twitter Login Join Twitter!

[@nagoodman](#) Everyone I know who's run Vectorwise tells me it's the fastest database they've seen.

11:47 AM Sep 17th via TweetDeck in reply to nagoodman
Retweeted by 1 person



markmadsen
Mark Madsen



twitter Login

THOLIS CONSULTING

Jos van Dongen
✉ jvdongen@tholis.com
☎ +31-651169606

I'm a consultant, author and presenter focussing on Open Source BI solutions. Recent book: 'Pentaho Solutions', published by Wiley: bit.ly/UyTVn

Website:
tholis.com

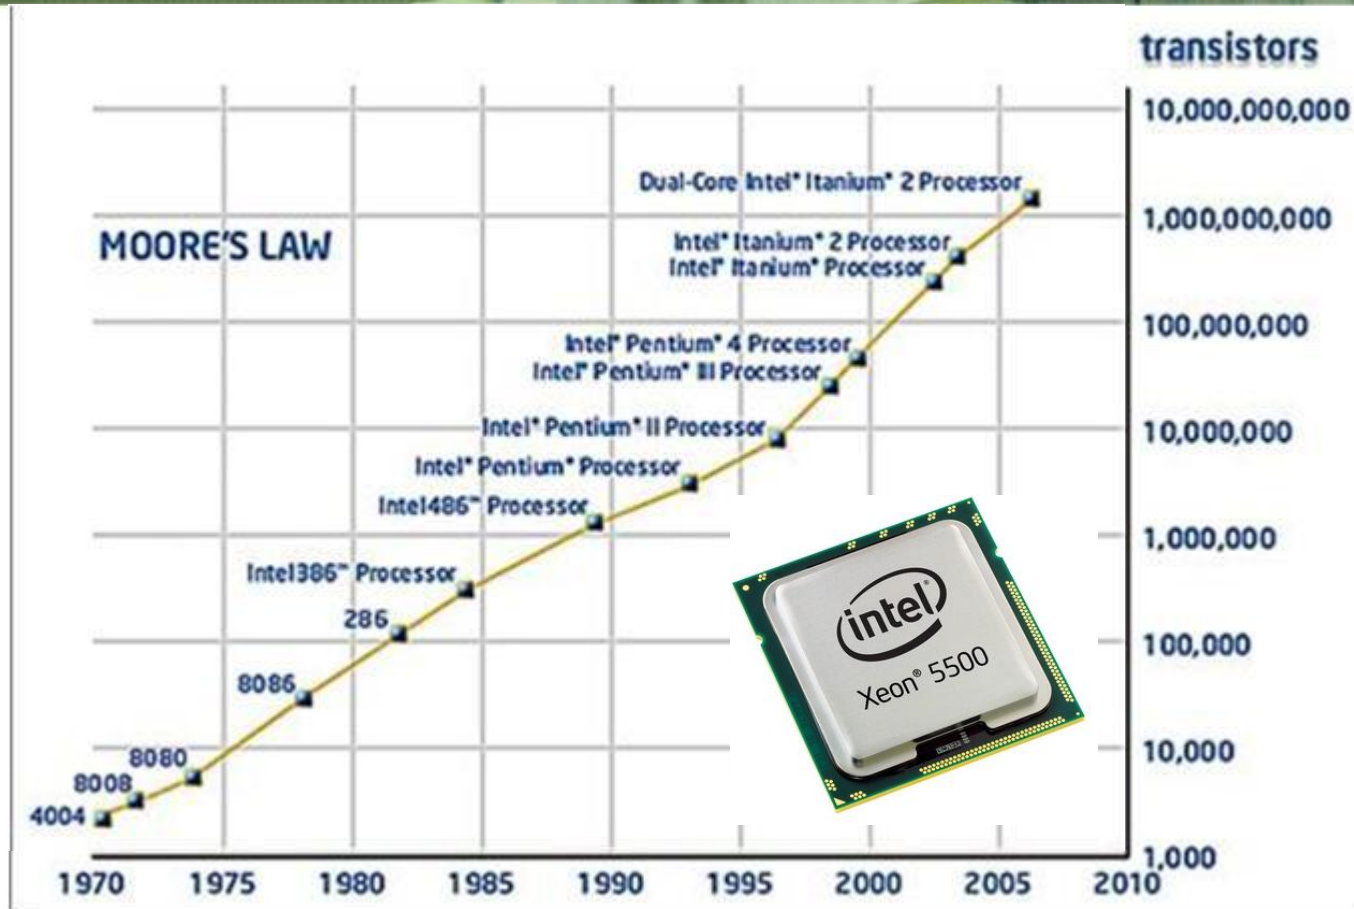
Comparing my own [@Vectorwise](#) results to published TPC-H SF100 numbers. Really amazing; maybe SAP bought the wrong engine ;-)

7:10 AM Jun 11th via HootSuite
Retweeted by 3 people



josvandongen
Jos van Dongen

Is the power of your data warehouse (/mart) increasing like this? This is Moore's law:



VectorWise unlocks the amazing potential of common chips

Performance of Data Warehouses and Marts

- **Performance is not everything - sure**
 - but it is an important factor in user satisfaction
 - if you achieved it, you may question at what financial, time and process cost you finally got it
- **Time-to-Market: from idea to production**
 - Do you have budget, IT support, implementation time?
 - **Simple** solutions are to be preferred!!
 - Cheap, simple to roll out, adapt to new requirements
 - Can be maintained without or with limited IT support
 - Performance is no longer this issue; focus on the business

Your Next Analytical DB. What do you prefer?

Traditional Warehouse

Complex, Custom and Expensive



Ingres VectorWise




(example machine is HP DL380, HP's entry-level server)

INGRES

vectorwise

INGRES

Resources to Manage Data Explosion are Constraint



Resources to manage data can't grow at same rate as data...

Physical data center

Human Resources

Budget

Energy cost

Expertise

Impending regulation

Businesses require a technology breakthrough to manage the data explosion

Our Key messages

- **Performance through breakthrough technology**
- **Simple, reduces cost and complexity**
- **Proven.**

History of VectorWise

- Centrum voor Wiskunde en Informatica, Amsterdam
 - Academic research institute
 - Active spin-off policy
 - 1999-2004 Data Distilleries BV (→SPSS→IBM)
 - 2006 ICT Regie Award
 - 2008 spin-off of VectorWise, with backing by Ingres Corp.



Database Innovation at CWI

- Pioneering column-stores
 - Architecture of choice OLAP/Warehousing/BI
 - MonetDB: 1994-
RAM-based open source column store
- Pioneering optimizing for CPU, not only I/O
 - Databases optimized for CPU Cache
 - Databases that exploit SIMD (SSE)

Academic Recognition:
2009 VLDB 10-year Award



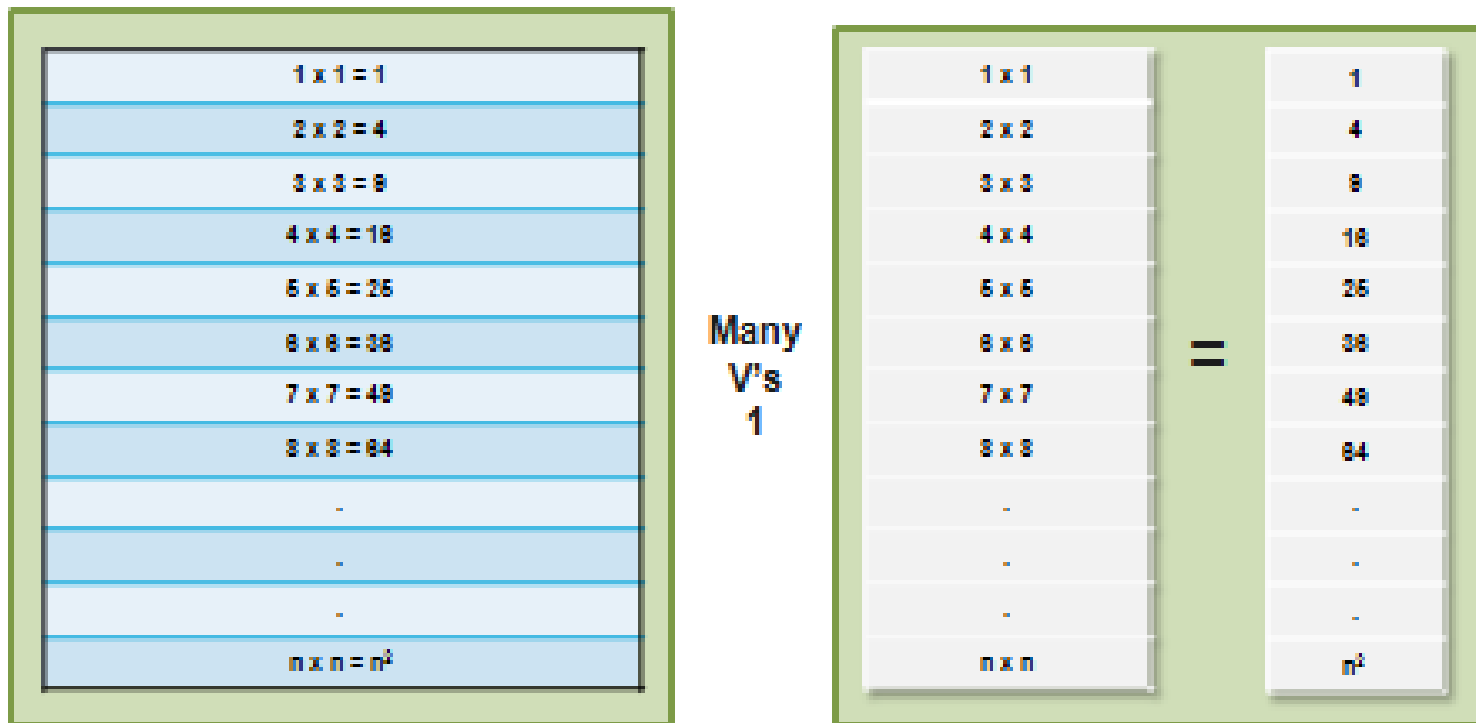
IBM Fellow Mohan (inventor of transaction systems) handed out the prize

VectorWise: Technical Highlights

- Secret Sauce: “vectorized execution”
 - All operations happen on “vectors” not on records
 - Makes the system 10-100x faster
- Innovation in compressed storage
 - Compression between RAM-CPU makes system less I/O dependent (less expensive configurations)
 - Super-fast vectorized algorithms
- Automatic Indexing and Clustering
 - Very good query plans without need of a specialized DBA
- Mixed Workloads (Operational BI)
 - Unique column-store update technology


Ingres VectorWise: Vector-based Processing

- **Single instruction performed on a set of data**
 - Vectorized computations, but also vectorized selection, vectorized sorting, vectorized joins



Ingres VectorWise: Processing in Chip Cache

- **VectorWise not only optimizes Disk but also RAM Access**
 - Operate on data that is in the CPU cache
 - Access to CPU cache orders of magnitude faster than RAM



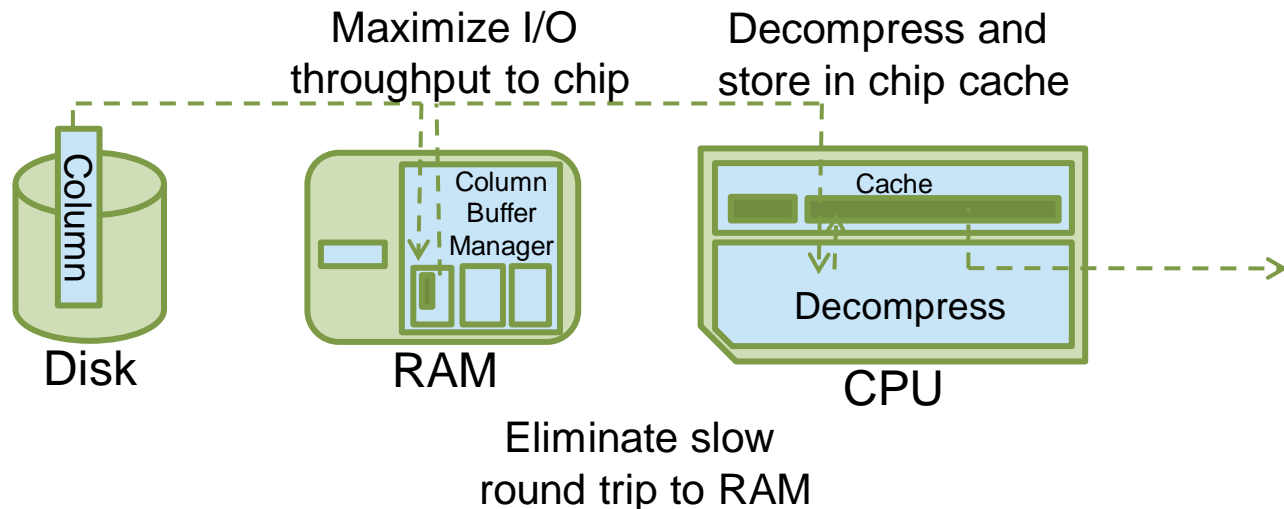
Ingres VectorWise: Updateable Column Store

- Only access relevant data, faster queries
- Enable incremental updates efficiently
 - Using new differential structures
 - Updates a weakness for other column-based DBMS

Cust_Num	Cust_surn ame	Cust_first_na me	Cust_mid_na me	Cust_DOB	Cust_Sex	Cust_Addr_1	Cust_Addr_2	Cust_City	Cust_State
46328927956	Jones	Steven	Sean	17-JAN-1971	M	333 SIGIda Rd		Melbourne	Vic
98679975745	Smith	Leonard	Patrick	04-APR-1964	M	Unit 12, 147 Trafalgar Sqr		Birmingham	London
52634346735	Rogers	Gindy	Carmine	11-MAR-1980	F	Belmont Rail Service	421 Station St	Belmont	CA
346737347347	Andrews	Jenny		14-SEP-1977	F	Apt1, 117 West 42nd St		New York	NY
88673477347	Cooper	Sheldon	Michael	30-JUN-1980	M	Ingres Corporation	Level 2, 426 Argello St	Redwood City	CA
34673447568	Koltwitz	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	Shanghai
-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-
53697659672	Cooperpoll y	Rahj	Henry	21-JUL-1978	M	Columbia State Library	17 Rodeo Dve	Los Angeles	CA

Ingres VectorWise: Compression

- **Extremely fast decompression between RAM and CPU**
 - Unique, new, fast compression methods
 - Automatically determined by Ingres VectorWise
- **Vectorised decompression**
 - Only for data processing in CPU cache



VectorWise: why Ingres?

- Commitment to make a major step
 - Dominant relational vendors not ready to throw out their existing Warehousing technology
- Established, stable, world-wide
 - Start-ups have to focus on few very big customers, with Ingres we have a global sales presence
- Small company focusing on Open Innovation
 - Ingres started a second life as an open source company with a flexible support business model
 - Ingres fosters and stimulates community developments (for VectorWise, there is a source licensing program)

Ingres Enterprise Support Organization



Proven!

Ingres has 10,000 + clients and business critical



ImPressing Solutions



BBP

Seamless connections



Lufthansa



the global loyalty agency



What is Ingres VectorWise

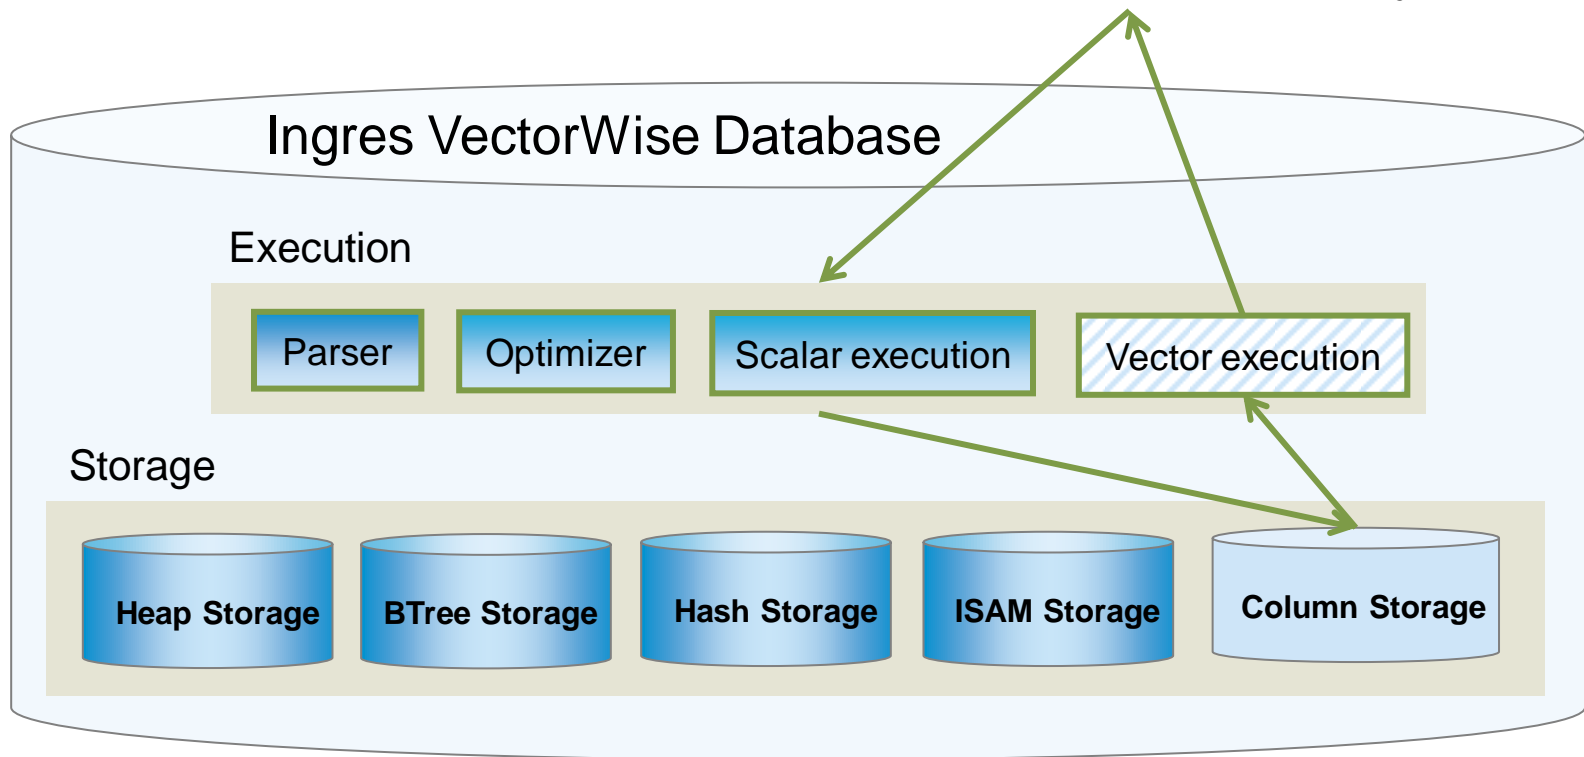
- **New column-based storage type in Ingres Database**
- **New execution engine in the Ingres Database, that**
 - Processes analytic queries much faster than any other RDBMS
 - Exploits performance features in modern CPUs
 - Uses Ingres Database to access and manipulate data

What is Ingres VectorWise

Existing ODBC/JDBC/.Net Drivers

ISO standard SQL

```
select ship_date, sum(quantity),...
```



Runs on commodity hardware

Ingres VectorWise Product Overview

- **Performance – through breakthrough technology**
 - Vectorised processing in the CPU's cache
 - Updateable column store, compression, storage index
- **Simple – reduce cost and complexity**
 - Transparent vectorised processing
 - Automatic compression and indexing
- **Proven**
 - Ingres Database and 24x7 follow-the-sun support
 - Technologies developed by CWI now being embraced by industry

Users have reported
10x – 75x performance gains
*using Ingres VectorWise
over other databases*

Beta Customers Experiences – Performance

*“For the past 20 years, we’ve been searching for the killer database that would fulfill our most intense data processing needs and with the discovery of Ingres VectorWise, that search is now over – this database technology is in **a class of its own**. Right out of the box, Ingres VectorWise lets us effortlessly **plow through millions and millions of rows of data**...The Ingres and VectorWise leaders and technologists have **performed a miracle here**.”*

Warren S. Master, CTO

The Rohatyn Group, a New York-based asset manager

*“We ported a business application from Oracle to Ingres VectorWise and were astounded by the **substantial performance increases, which were up to 70x**...With Ingres VectorWise Database, we get the best of both worlds; we can now offer our customers a lower cost and better performing analytic database.”*

Michael Thuleweit, Managing Director, Datamatics, Frankfurt

Beta Customers Experiences – Simplicity

*“The speed **is blindingly fast right out of the box** and **it eliminates layers of work that had to be done in the past**. Ingres VectorWise uses automated compression and indexing which **eliminates the need to spend extra hours tuning and tweaking our data**. This means we are now able to cut the time of our business intelligent engagements with our customers by up to 50%.”*

Roy Hann at Rational Commerce, a leading systems integrator

*“Ingres VectorWise performance has been outstanding and it was so **easy to get up and running**. Position analysis was **10 to 50 times** faster than the leading commercial databases in our testing and with linear query performance from 12 mm – 1.2 billion financial records. **These are drastic improvements in terms of speed.**”*

Steve Ferrando, CEO, dbConcert, a leading systems integrator to major Wall Street investment firms

The Value of Simple

■ Competitors

- Main relational vendors
 - OLAP performance may lack, and is very tuning intensive
- Cluster systems:
 - take eg PostgreSQL (slowwww) and put many on a cluster
 - Exaggerated hardware cost, hardware maintenance = DBA time (replacing disks, maintaining nodes, etc)
- Hardware solutions (e.g. Nettezza with FPGA)
 - costly black box, maximal lock-in

THINK IT
VECTORISE IT
DONE

■ Ingres VectorWise

- A warehouse that performs like a cluster but is an easy-to-administer single node system (cheap, simple and green)
- Thanks to unique and acclaimed Dutch technology!
- Flexible enterprise support subscription model



Questions?

- **[Visit: ingres.com/vectorwise](https://ingres.com/vectorwise)**

Gerard.Ringenaldus@ingres.com