IBM SolutionsConnect 2013
Turning Opportunity into Outcomes.

The Next BIG Thing - BLU

Eric Thijs

Agenda

1) The Next BIG Thing – BLU acceleration by Eric Thijs, IBM
    1) General
    2) DB2 – Doing More … with Less …
        1) Performance
        2) Availability
        3) Modern Technology
        4) SW costs
        5) HW investments
        6) Operational effort
        7) Migration, Support and Education
        8) Question Round
2) Customer Story by **Mr Reinoud Reynders**, IT-Manager Infrastructure & Operations @
    1) UZ Leuven – Migration from Sybase to DB2
        1) Goals & Challenges
        2) Results
        3) Lessons Learned
        4) Question Round

General

Factors Impacting Organisations

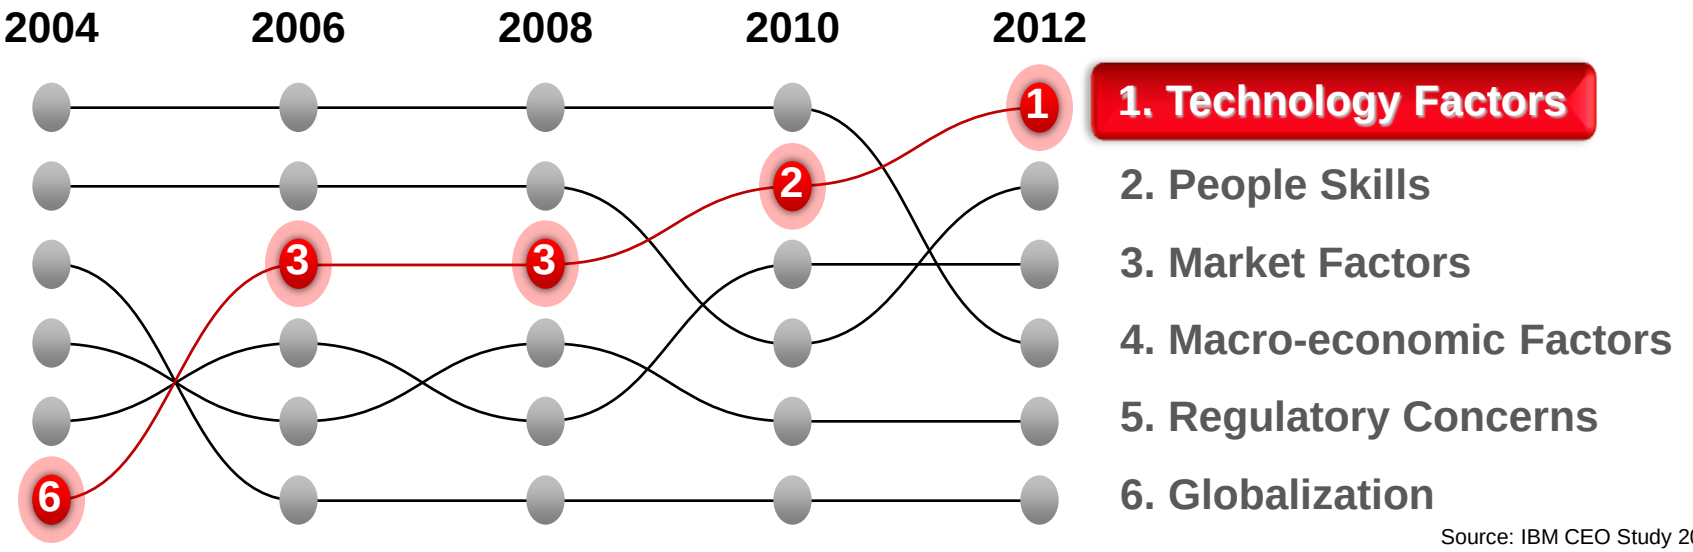

Source: IBM CEO Study 2012

2.7ZB

Digital content

1.2B

Mobile employees

556M

Cybercrime victims

3.6X

Competing on analytics

General

IT Architecture Complexities Stifle Insight and Action

General

IT Infrastructure under tremendous pressure

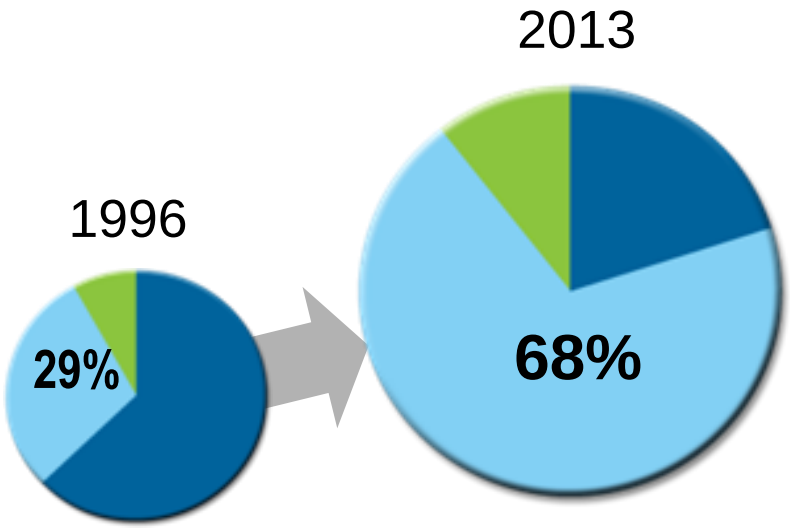

68% of IT

operating costs in 2013 will be for **management and administration**

* IDC; Converged Systems: End-User Survey Results presentation; September 2012; Doc #236966

Only 1 in 5

organizations allocate more than **50% of IT budget to new projects**

* 2012 IBM Global Data Center Study

General

IBM Database overview

1. DB2 – Hybrid DB
    1. SAP certified
    2. Non-SAP

2. Informix – Extreme stability and massive data congestion
    1. Time series – Smart meters & Sensor Data
    2. Flexible GRID – Localisation
    3. Real time analytics with IWA (Informix Warehouse Accelerator)

3. SolidDB – In-Memory database for millisecond response
    1. Telecom & Banking
    2. SolidDB mobile

DB2
"Doing More with Less"

Doing MORE …

- Performance
  - Hybrid DB (OLTP & OLAP)
  - Lesser I/O and better HW utilization
  - DPF, MDC, NLS, …
- Availability
  - No more planned downtime
  - Stability (lesser patches, 10 x)
  - Reduced Backup / Restore time
- Modern Technology
  - SAP partnership
    - Joint roadmap
    - Shortest certification track
  - BLU acceleration
- Integration - SAP

… With LESS

- Software Cost
  - Lower License and Maintenance
  - HADR, DPF, pureScale, BLU, …
- Hardware Investments
  - Storage
    - Disks and tapes
  - Memory
  - Servers
- Operational effort, staff efficiency
  - Automation
    - ASM, STMM, WLM, …
  - Tooling
    - Workload Manager, DBA cockpit

DB2

More Performance

- Present DB2 10.1
  - OLTP – minimal 20% performance gain
  - OLAP – minimal 40% performance gain x
- New DB2 10.5 - OLAP
  - Columnar Stored (5) – Slide 10
  - Optimal Memory caching (6) – Slide 11
  - Data Skipping (7) – Slide 12

Customer	Speedup over DB2 10.1
Large Financial Services Company	46.8x
Global ISV Mart Workload	37.4x
Analytics Reporting Vendor	13.0x
Global Retailer	6.1x
Large European Bank	5.6x

10x-25x
improvement
is common

DB2
More Availability

- Present DB2 10.1
  - Robust and stable technology
    - Patch level
    - Support
  - HADR – High Availability and Disaster Recovery
    - Primary & up to 3 auxiliary sites
    - Rolling fixpack upgrades – no more planned downtime
  - pureScale – Continuous Availability and Unlimited Scalability
- New DB2 10.5
  - Next Fixpack/Release
    - HADR & pureScale
    - More than AIX / Power HW

DB2
More Modern Technology

- Present DB2 10.1

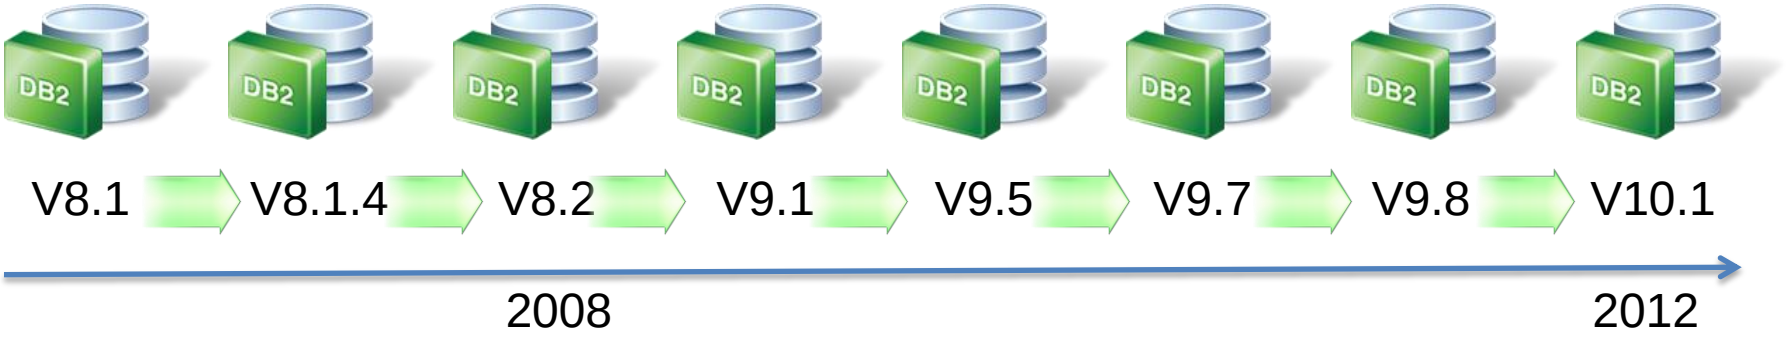

- New DB2 10.5 - 2013
  - BLU (Blink Update) acceleration
  - 7 BIG ideas
    - Create & Load (1) – convert or create
    - Compression (2)
    - SIMD (3)
    - Core Friendly Parallelism (4)
    - Columnar stored (5)
    - Scan Friendly Memory Caching (6)
    - Data Skipping (7)

DB2

Lesser software costs

- Present DB2 10.1
  - SAP model
    - 8% of the SAV (Hardware independent model)
    - ESE + SOF + DPF + HADR + BLU
  - Non-SAP
    - lesser PVU's
    - Flexible licensing models
- New DB2 10.5
  - SAP model
    - Included in the 8% of the SAV
  - Non-SAP
    - Developed to run on existing HW
    - Simplified Licencing model x

DB2
Lesser hardware investments

- Present DB2 10.1
  - SOF (storage Optimization Feature – Adaptive Compression
    - Lesser CPU, RAM and Storage
    - Less I/O

- New DB2 10.5 – BLU ac
  - Higher compre
  - SIMD – Single
  - Core Friendly F
  - Columnar Stor
  - Scan friendly n

Slide 11

Table Courtesy of Triton, Iqbal Goralwalla	Size (GB)	Storage Savings
T1 Row organized	185	N/A
T2 – T1 Adaptive C	57,4	69% or 3,23x
T3 – Convert T1 to column with db2convert	12	93,5% or 15,4x versus T1 79,1 or 4,78xversus T2
T4 – Create new column organized table like T1 and load data	12	Same as T3

DB2
Lesser operational effort

- Present DB2 10.1
  - Automations
    - STMM, ASM, …
    - Between 20% and 25% lesser effort to maintain
  - Integrated tooling
    - SAP – DBA cockpit
    - Non-SAP – Optim Performance Manager (AESE)
- New DB2 10.5
  - Simplification – “Create & Load” (1) – Slide 24
    - Built seamless into the DB2 kernel
      - Integration and
        - Same SQL,
        - Language interfaces &
        - Administration
      - Coexistence - Column-organized tables can coexist with existing traditional tables
        - Same schema,
        - Storage &
        - Memory
  - Compression (2) see also HW investments – Slide 20

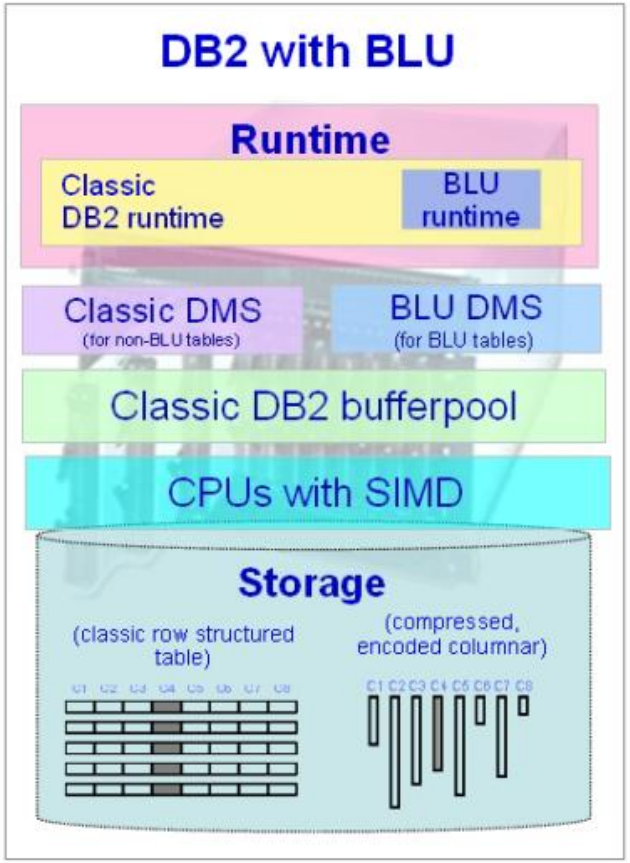

DB2

Migrations, Support and Education

- SAP
  - FIXED Migration Offer – Migration Sizing Questionnaire
    - T-shirt size – based on 1 landscape with 3 instances
      - Small - Production Database > 500 GB - € 15K
      - Medium - Production Database < 500 GB > 1 TB - € 20K
      - Large- Production Database < 1 TB - € 25K
  - 100% success ratio @ 1.000 customer and 10.000 instances
    - Certified Tooling, People and processes
- Non-SAP
  - Oracle compatibility 98%
  - FREE assessment and estimated conversion report

- Support
  - Workshops – Deep Dives
  - One Day Migration planning

- Education
  - Courses
    - Classroom
    - Webcasts

DB2
Question round

Thank You

Eric Thijs

Information Management | SWG

Core DB | Infosphere Guardium |
PureData for Transactions Leader BeLux

+32 478 630 473

eric.thijs@be.ibm.com

DB2

More Performance – Performance results 9.7 & 10.1

Customer	Application	DB version	Source (H:min:sec)	DB2 after	Reduction	Faster %	
Swiss P&A	40 SAP Jobs	11.2	23:58:48	13:33:27	10:25:21	43,46%	
	6 Batch		20:56:48	10:43:27	10:13:21	48,80%	
Audi	ERP	10.2	14:43:44	5:03:27	9:40:17	65,66%	40,60%
	BW		3:36:44	3:18:27	0:18:17	8,44%	
	eRecruiting		0:13:20	0:08:27	0:04:53	36,63%	
Belgian comp	OLTP	10.2	3:12:12	1:04:36	2:07:36	66,39%	
	OLTP		1:20:44	0:58:27	0:22:17	27,60%	
	Batch		0:00:40	0:00:27	0:00:13	32,50%	
ABB	Online perf.		0:48:40	0:10:42	0:37:58	78,01%	
	Batch perf.		1:18:00	0:39:27	0:38:33	49,42%	54,97%
	Online perf.	11.2	0:34:04	0:10:42	0:23:22	68,59%	
	Batch perf.		2:43:48	0:39:27	2:04:21	75,92%	
Swiss P&A	40 SAP Jobs		23:58:48	12:59:27	10:59:21	45,83%	
	6 Batch		20:56:48	10:22:27	10:34:21	50,47%	

DB2

More Performance – Performance results BLU by Triton

Table Courtesy of Triton, Iqbal Goralwalla	Sample Workload Total response time (Sec)	Sample Workload Response time speedup
T1 – Row organized	1.385	N/A
T3	31	45x
T4	31	Same as T3

Coca-Cola Bottling Co. Consolidated

"The BLU Acceleration technology has some obvious benefits: It makes our analytical **queries run 4-15x faster** and **decreases the size of our tables by a factor of 10x**. But it's when I think about **all the things I don't have to do with BLU**, it made me appreciate the technology even more: **no tuning, no partitioning, no indexes, no aggregates**."

-Andrew Juarez, Lead SAP Basis and DBA

DB2

More Performance – Columnar Store (5)

- Minimal I/O
  - Only perform I/O on the columns and values that match query
  - As queries progresses through a pipeline the working set of pages is reduced
- Work performed directly on columns
  - Predicates, joins, scans, etc. all work on individual columns
  - Rows are not materialized until absolutely necessary to build result set
- Improved memory density
  - Columnar data kept compressed in memory
- Extreme compression
  - Packing more data values into very small amount of memory or disk
- Cache efficiency
  - Data packed into cache friendly structures x

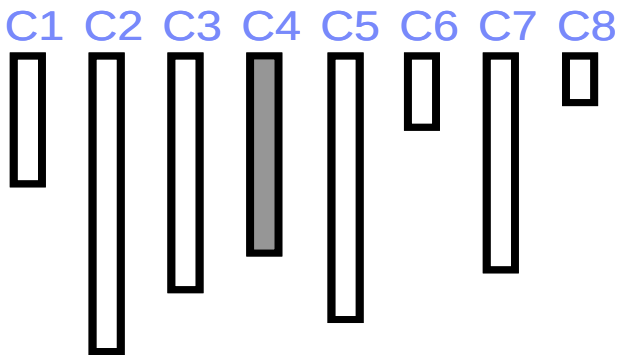

DB2

More Performance – Optimal Memory Caching (6)

- New algorithms cache in RAM effectively
- High percent of interesting data fits in memory
  - We leave the interesting data in memory with the new algorithms
- Data can be larger than RAM
  - No need to ensure all data fits in memory
  - Optimization for in memory and I/O efficiency x

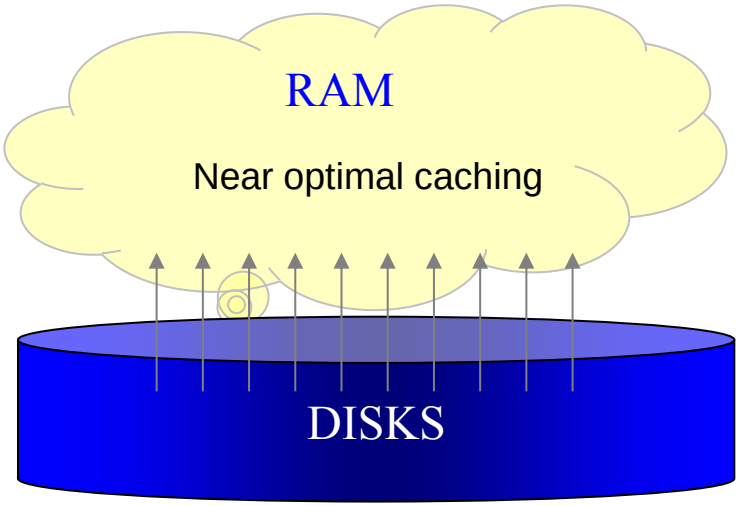

DB2

More Performance – Data Skipping (7)

- Automatic detection of large sections that do not qualify for a query and can be ignored
- Order of magnitude savings in all of I/O, RAM, and CPU
- No DBA action to define or use – truly invisible
  - Persistent storage of min and max values for sections of data values x

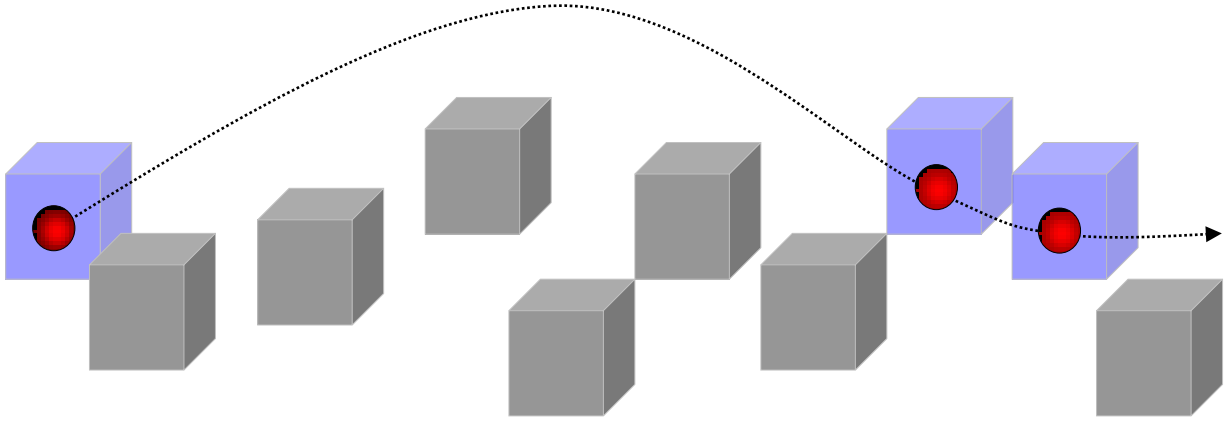

DB2

Lesser software costs – Simplified Licence model

	Departmental Market	Enterprise Market
Advanced function	**DB2 Advanced Workgroup Server Edition** • For small OLTP and analytic deployments • Primarily used in department environments within large enterprises or SMB/MM deployments • Limited by TB, memory, sockets, and cores. • Supports BLU, pureScale and DPF deployment models	**DB2 Advanced Enterprise Server Edition** • For enterprise class OLTP and/or analytic deployments. • Targeting full enterprise/full data centre requirements • No TB, memory, socket, or core limit • Supports BLU, pureScale and DPF deployment models
Core function	**DB2 Workgroup Server Edition** • Entry level offering • Single server for less intense workloads • Limited by TB, memory, sockets, and cores • No support for BLU, pureScale or DPF deployment models	**DB2 Enterprise Server Edition** • Entry level offering. • Single server for enterprise/more intense workloads • No TB, memory, socket, or core limit • No support for BLU, pureScale or DPF deployment models

Limited capacity	*Full capacity*
DB2 Developer Edition	**DB2 CEO**
DB2 Express and DB2 Express-C	**DB2 Advanced CEO**

DB2

Lesser hardware investments – Compression results 9.7 & 10.1

SAP System	Other DB in GB	DB2 9.7 and 10.1 in GB	Compression in %
SAP ERP	1.145,361	329,633	71.22%
SAP ERP	16.965,628	3.767,178	77.80%
SAP BW	6.460,350	1.990,323	69.19%
SAP ERP	1.922,977	452,041	76.49%
SAP ERP	1,797820	462,560	74.27%
R/3 4.7 Enterprise *	630,880	229,040	63.70%
SAP ERP	598,229	151,112	74.74%
SAP ERP	425,098	119,102	71.98%
SAP ERP	643,860	160,140	75.13%
R/3 4.7 Enterprise *	60,416	20,480	66.10%
SAP ERP	557,990	153,770	72.44%
SAP BW	4.527,104	1.363,968	69.87%
SAP ERP	3.474,432	911,360	73.77%
TOTAL	**39.210,145**	**10.110,707**	**74.21%**

* With DB2 9.7

DB2

Lesser hardware investments – Compression results BLU

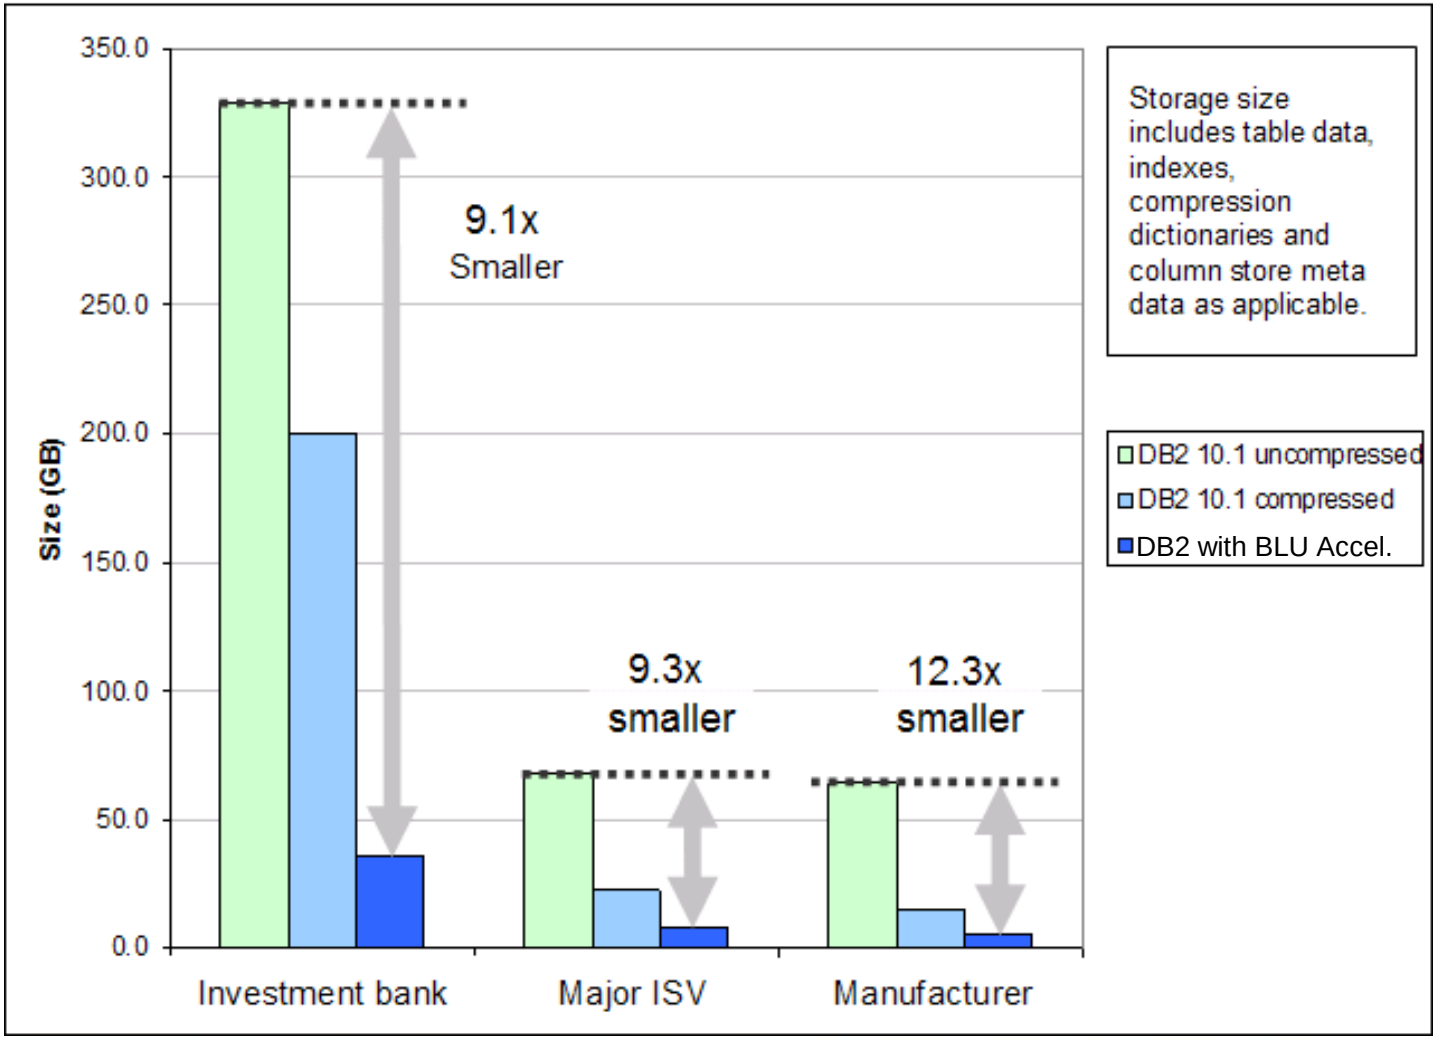

DB2

Lesser hardware investments – Compute Friendly & Compression

- Massive compression with approximate Huffman encoding
  - More frequent the value, the fewer bits it takes
- Register-friendly encoding dramatically improves efficiency
  - Encoded values packed into bits matching the register width of the CPU
  - Fewer I/O's, better memory utilization, fewer CPU cycles to process

LAST_NAME	Encoding
Johnson	
Smith	
Smith	
Smith	
Smith	
Johnson	
Smith	
Gilligan	
Sampson	
Smith	

Packed into register length

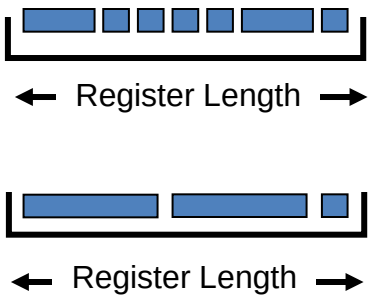

DB2

Lesser hardware investments – Data remains compressed

- Encoded values do not need to be decompressed during evaluation
  - Predicates (=, <, >, >=, <=, Between, etc), joins, aggregations and more work directly on encoded values

```
SELECT COUNT(*) FROM T1 WHERE LAST_NAME = 'Johnson'
```

LAST_NAME	Encoding
Brown	
Johnson	
Johnson	
Johnson	
Johnson	
Brown	
Johnson	
Gilligan	
Wong	
Johnson	

DB2

Lesser hardware investments – SIMD (3)

- Performance increase with Single Instruction Multiple Data (SIMD)
- Using hardware instructions, DB2 with BLU Acceleration can apply a single instruction to many data elements simultaneously
  - Predicate evaluation, joins, grouping, arithmetic

intel

*"Intel is excited to see a **25x improvement in query processing performance** using DB2 10.5 with BLU acceleration over DB2 10.1. To achieve these amazing gains, IBM has taken advantage of the Advanced Vector Extensions (AVX) instruction set on Intel® Xeon® processor E5-based systems."*

- Pauline Nist, GM, Enterprise Software Alliances, Datacenter & Connected Systems Group

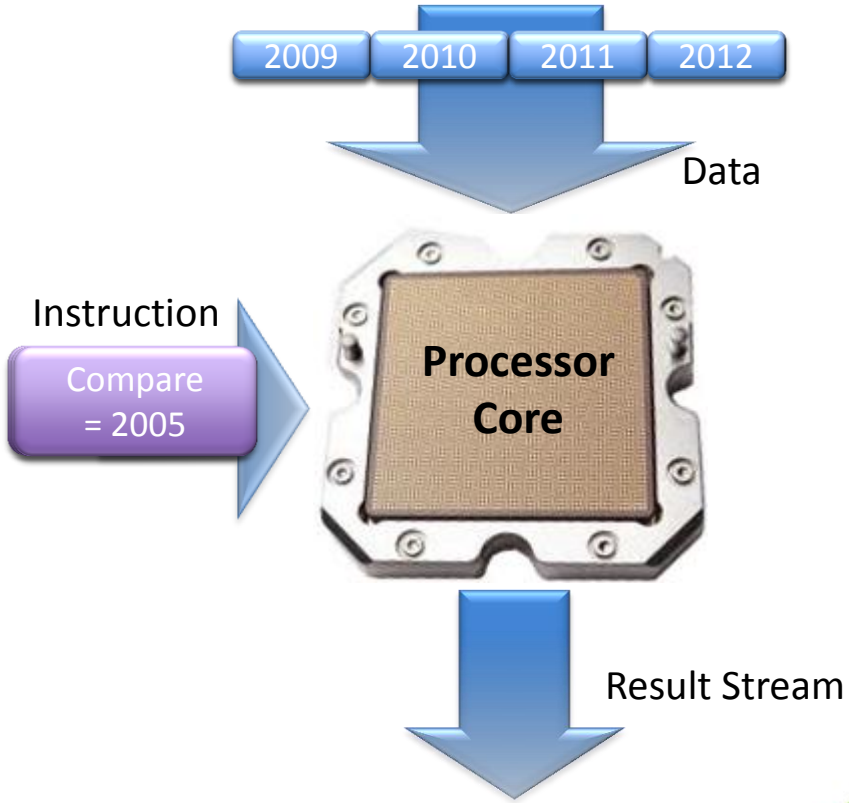

DB2

Lesser hardware investments – Core Friendly Parallelism (4)

- Careful attention to physical attributes of the server
  - Queries on BLU Acceleration tables automatically parallelized
- Maximizes CPU cache, cacheline efficiency

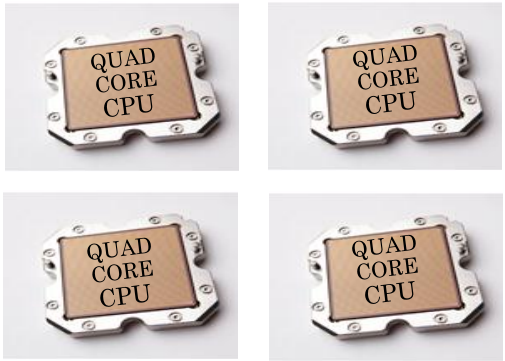

BNSF RAILWAY

"During our testing, we couldn't help but notice that ***DB2 10.5 with BLU Acceleration is excellent at utilizing our hardware resources****. The core-friendly parallelism that IBM talks about was clearly evident and* ***I didn't even have to partition the data across multiple servers****."*

- Kent Collins, Database Solutions Architect, BNSF Railway

DB2

Lesser operational effort – Create & Load

DATABASE

MICROSOFT
SYBASE
TERADATA
ORACLE

Database Design and Tuning

- Decide on partition strategies
- Select Compression Strategy
- Create Table
- Load data
- Create Auxiliary Performance Structures
  - Materialized views
  - Create indexes
    - B+ indexes
    - Bitmap indexes
- Tune memory
- Tune I/O
- Add Optimizer hints
- Statistics collection

Repeat

DB2
General Concept

32 cores
1TB memory
10TB table
100 columns
10 years data

```
SELECT COUNT(*) from MYTABLE
where YEAR = '2010'
```

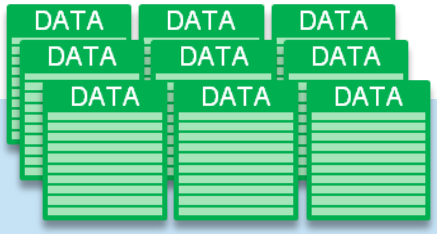

10TB data

Actionable Compression
reduces to **1TB**
In-memory

Column Processing
reduces to **10GB**

Data Skipping
reduces to **1GB**

Massive Parallel Processing
32MB linear scan
on each core

Vector Processing
Scans as fast as
8MB through **POWER7**
Accelerated SIMD

Result in seconds or less