


Breakthrough Analytics Performance with BLU Acceleration

Answers At The Speed Of Thought
To Help You Outperform The Competition



Businesses Of All Types Use Analytics To Accelerate Decision Making At The Point Of Impact

To predict and treat high-risk patients to proactively intervene at time of visit

To identity and prevent fraud in real time

To provide targeted cross sells that drive additional sales when offered at time of sale

To predict correct order levels at time of purchase

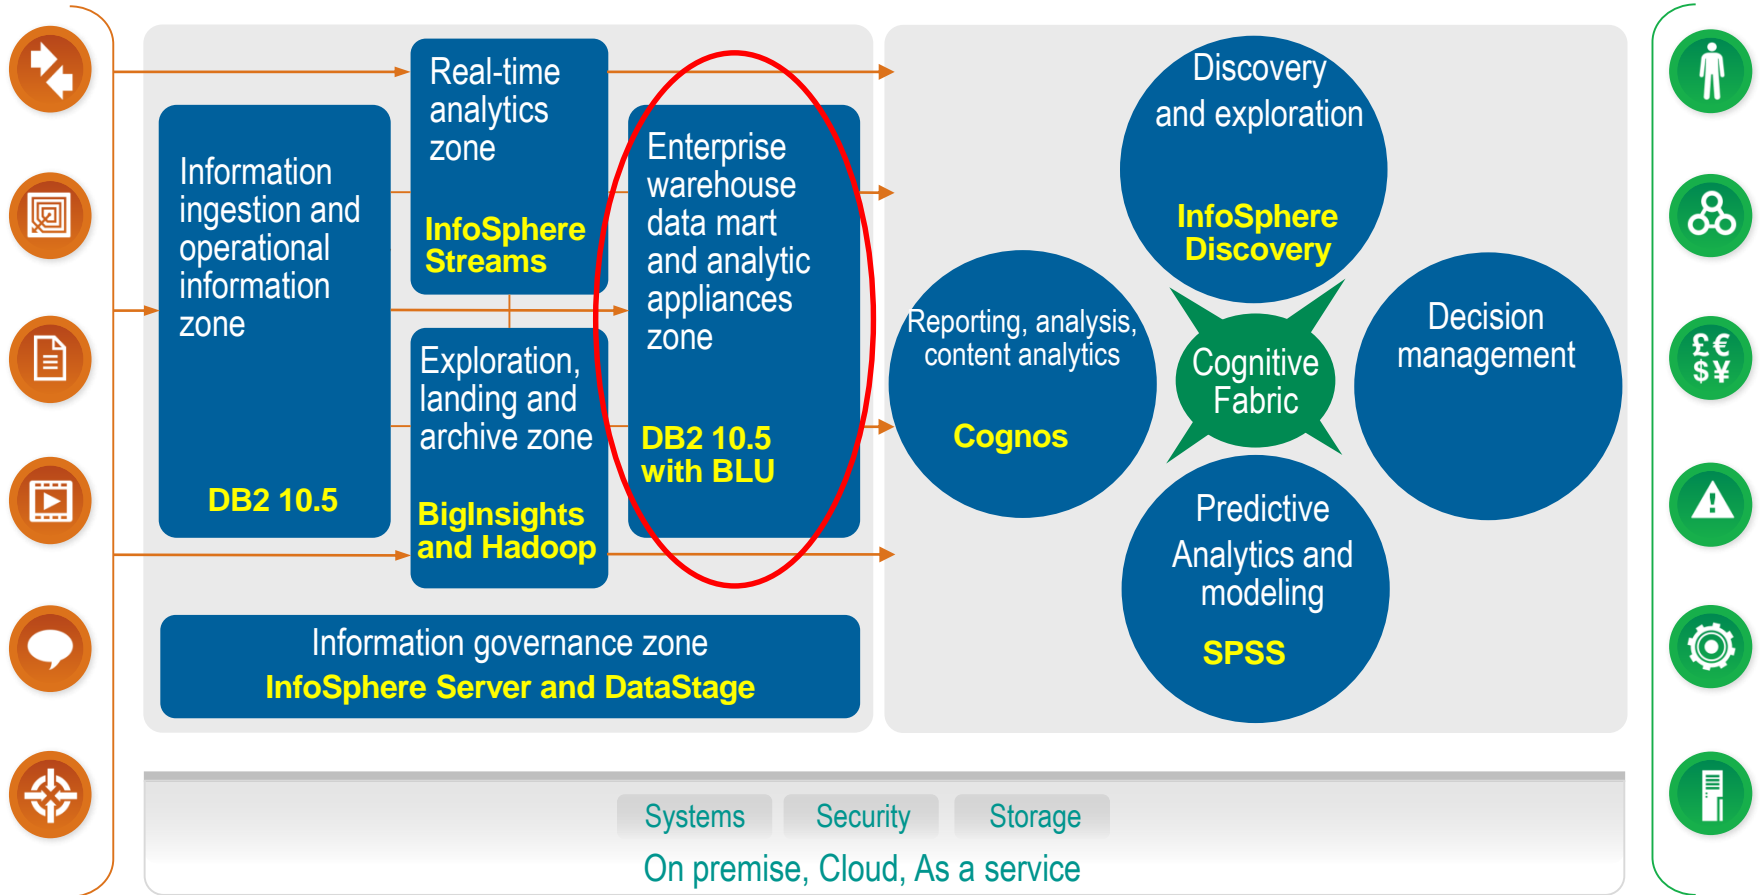


Make Sense out of BigData with IBM Watson Foundations

All Data

IBM Watson Foundations

New/Enhanced Applications



Analyze all data, from any source, with the right technology

DB2 with BLU Acceleration

Fast Answers. Simply Delivered.

What is DB2 with BLU Acceleration?

- In-memory analytic database
- Multiple IBM innovations
 - In-memory processing of columnar data without the limitations of memory size
 - Analyze compressed data with actionable compression
 - CPU Acceleration
- Ready for Analytics: Cloud, On premise, SAP, Cognos, and more
- Agile warehousing via BLU for Cloud
- www.ibmBLUhub.com




BLU Acceleration

*Analyze more data
faster and more
efficiently*

DB2 with BLU Acceleration

What more can businesses do with BLU?



How many more questions could your business answer if they could get answers this much faster?



Speed of Thought Analytics

- Analyze **data literally as fast as you can ask questions**
- “We’ve tested DB2 10.5 with BLU Acceleration and found that it can be up to **43x faster** with an analytic workload...” - Randy Wilson, Lead DB2 for LUW DBA

What can you do for your business with an extra month?



Operational Simplicity

- Maximizes business value from existing infrastructure; no need to rip and replace
- “We project this will save us 42 days per year in lower administration and tuning efforts” – Brenda Boshoff, Sr. DBA

Where could you invest the 10x savings to drive more business value?



Business Agility

- Deploy on-premise or via the cloud with only a fraction of the computing resources
- “Using DB2 10.5 with BLU Acceleration, our storage consumption went down by about **10x**” - Kent Collins, Database Solutions Architect



What Makes BLU Acceleration Different?

Unmatched Innovations From IBM Research & Development Labs

Next Generation In-Memory

In-memory columnar processing with dynamic movement of data from storage




Analyze Compressed Data

Patented compression technique that preserves order so data can be used without decompressing




CPU Acceleration

Multi-core and SIMD parallelism (Single Instruction Multiple Data)



Data Skipping

Skips unnecessary processing of irrelevant data




First, let's compare *column* organization of data with traditional *row* organization of data

Row Organized Data is Well Suited and Efficient for Transactional Workloads (OLTP)

- Traditional systems typically organize data into rows or “records”
- One or more rows are combined into “pages” and stored on disk
- A record contains a value for each column, stored together on the same page
- One I/O operation (to disk or RAM) can retrieve all column values for a record
- Great for transactional workloads
 - When MANY columns are accessed for SPECIFIC records
 - Indexes are required to improve performance when accessing specific records

Row Organized Customer Table

	CUST_ID	FIRST	LAST	AGE	SEX
Row 1	466	Steve	Miller	49	M
Row 2	467	Pat	Smith	32	F
Row 3	478	Tina	Jones	27	F
Row...	479	Rick	Miller	42	M
Row N	481	Tom	Smith	36	M




Query:

```
Select FIRST, LAST, AGE, SEX
from Customer where CUST_ID=466
```

466	Steve	Miller	49	M
-----	-------	--------	----	---

Efficient!

Query needed 4 columns, I/O retrieved the entire page, the entire row, and all the required column values for a record.

 Each colored row represents a data page

Row Organized Data Can be Inefficient for Analytic Workloads

- Analytics queries often operate on only a small number or even a single column value across a very large number of rows
 - ▶ For example: MIN, MAX, SUM, COUNT, AVG
- Retrieving all column values is inefficient when only a small number of columns (maybe just 1) are needed

Row Organized Customer Table

	CUST_ID	FIRST	LAST	AGE	SEX
Row 1	466	Steve	Miller	49	M
Row 2	467	Pat	Smith	32	F
Row 3	478	Tina	Jones	27	F
Row...	479	Rick	Miller	42	M
Row N	481	Tom	Smith	36	M

■
■
■
■
■
 Each colored row represents a data page



Query:

Select AVG(AGE) from Customer

466	Steve	Miller	49	M
467	Pat	Smith	32	F
478	Tina	Jones	27	F
479	Rick	Miller	42	M
481	Tom	Smith	36	M

Less Efficient! **AVG=37.2**

Query only needed 1 column to compute average age, but I/Os were required to retrieve all columns for all rows

Column Organized Data is Better Suited and More Efficient for Analytic Workloads

- BLU Acceleration organizes data into columns
- Column values for many records are combined into “pages” and stored on disk
- One I/O operation (to disk or RAM) can retrieve a column value for many rows
- Great for analytical workloads
 - ▶ When SPECIFIC columns are accessed for MANY records
 - ▶ No indexes required – columns are essentially “self indexing”

Column Organized Customer Table

CUST_ID	FIRST	LAST	AGE	SEX
Col A	Col B	Col C	Col ...	Col N
466	Steve	Miller	49	M
467	Pat	Smith	32	F
478	Tina	Jones	27	F
479	Rick	Miller	42	M
481	Tom	Smith	36	M

■
■
■
■
■
 Each colored column represents a data page

Query:

Select AVG(AGE) from Customer



49
32
27
42
36

Efficient!

AVG=37.2

Query needed 1 column to compute average age, I/O retrieved just one column

BLU Acceleration is Integrated into DB2 10.5

BLU Acceleration



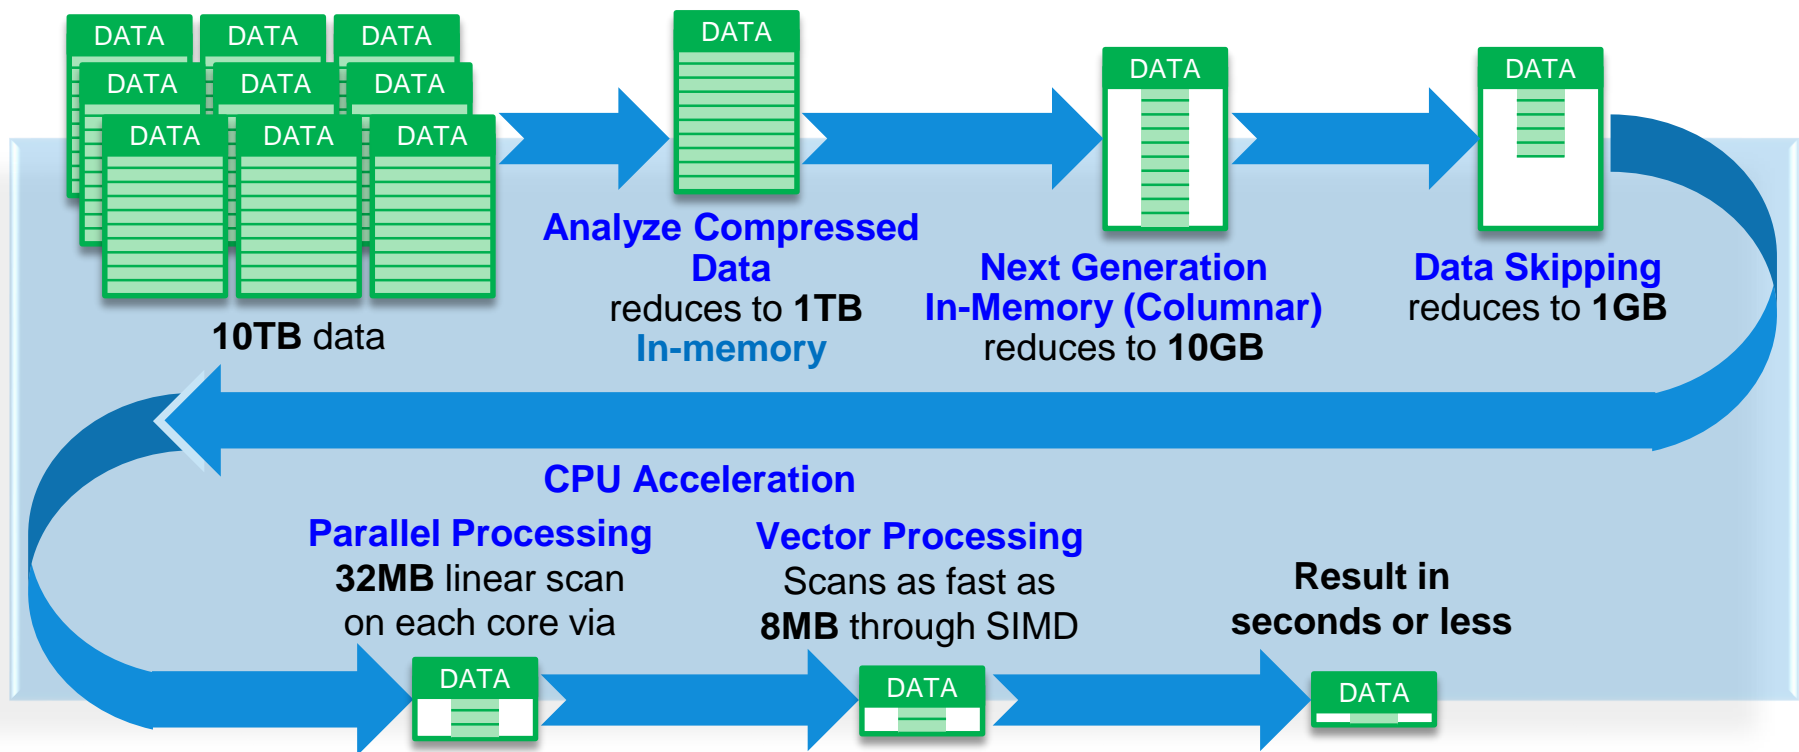
- One setting to optimize the system for BLU Acceleration
- Automatically configures DB2 for optimal analytics performance
- Easily convert tables from row-organized to column-organized
- Simple to create BLU tables
 - “organize by column”

DB2 `WORKLOAD=ANALYTICS`

- Makes column-organized tables the default table type
- Enables automatic workload management
- Enables automatic space reclaim
- Page and extent size configured for analytics
- Memory for caching, sorting and hashing, utilities are automatically initialized based on the server size and available RAM
- Compression is always on

DB2 BLU Acceleration Illustration: 10TB query in Seconds or Less

- The System: **32 cores, 1TB memory, 10TB table with 100 columns and 10 years of data**
- The Query: **What were our average, total, and number of “sales” in 2010 by month?**
 - `SELECT MONTH, AVG(SALES), SUM(SALES), COUNT(*) from MYTABLE where YEAR = '2010' group by MONTH`
- The Result: **In seconds or less as each CPU core examines the equivalent of just 8MB of data**



DEMO: DB2 10.5 with BLU Acceleration

- Two fact tables each loaded with 500M records
 - ▶ Uncompressed data size = 55GB
 - ▶ BLU table, 9.6GB compressed (5.7x), 5GB buffer pool
 - ▶ Row-organized table, 14.52GB compressed (3.8x), 12GB buffer pool
- **Compare** performance of BLU Acceleration table vs. traditional row-organized table

<i>Query Description</i>	BLU Acceleration Advantage
Query 1 Count the total number of records in the fact table (500 million)	14 X
Query 2 Calculate the average profit per sale for all 500 million records	8 X

BLU Acceleration Allows Super Fast Analytics While Simplifying Deployment and Maintenance

- BLU Acceleration is SIMPLE and EASY to use!
 - BLU column-organized tables provide great analytics performance while making things simpler and easier for DBAs
 - BLU Acceleration provides great performance, no tuning required
 - Alleviates the need for traditional warehouse performance structures
 - Administrators spend more time on application capability

- LOAD and then... run queries
 - No indexes
 - No REORG (it's automated)
 - No RUNSTATS (it's automated)
 - No Materialized Query Tables or materialized views
 - No partitioning or Multi-Dimensional Cluster indexes
 - No optimizer hints

BLU Acceleration Beats the Competition

What about SAP HANA and Oracle Exadata? They've made bold performance claims...



**Service Oriented Finance
CIO**

BLU Acceleration is better, faster, simpler, and costs less. We'll show you test results to prove it.



IBM

The BI Day Serial Execution Test Measures Time for One User to Complete Individual Reports

1 user executes **complex** reports ...then... executes **intermediate** reports ...then...executes **simple** reports




1 User



Each report executes one or more queries

Single User Connection



- SQL generated by Cognos
- Single user test, serial execution
- Each query gets all system resources

Examples of Simple and Complex Queries

Example of a typical **simple** query: 2-3 parameters compared

```
- QUERY 2A Simple_GoBusinessView_Dashboard select distinct "Product_forecast"."YEAR" "Year1" from "GOSL"."PRODUCT_FORECAST"
"Product_forecast"
```

Example of a typical **complex** query: 100-150 parameters compared


```
-- QUERY 3A Complex_TPA_PDF
with "Order_method_dimension14" as (
select "Order_method_dimension"."ORDER_METHOD_KEY" "ORDER_METHOD_KEY", min("Order_method_dimension"."ORDER_METHOD_EN") "Order_method"
from "GOSLDW"."ORDER_METHOD_DIMENSION" "Order_method_dimension"
group by "Order_method_dimension"."ORDER_METHOD_KEY"), "Product_line15" as (
select "Product_line"."PRODUCT_LINE_CODE" "PRODUCT_LINE_CODE", min("Product_line"."PRODUCT_LINE_EN") "Product_line"
from "GOSLDW"."PRODUCT_LINE" "Product_line"
group by "Product_line"."PRODUCT_LINE_CODE"), "Sales_territory_dimension12" as (
select "Sales_territory_dimension"."COUNTRY_KEY" "COUNTRY_KEY", "Sales_territory_dimension"."COUNTRY_CODE" "COUNTRY_CODE",
"Sales_territory_dimension"."SALES_TERRITORY_KEY" "SALES_TERRITORY_KEY", "Sales_territory_dimension"."SALES_TERRITORY_CODE" "SALES_TERRITORY_CODE",
"Sales_territory_dimension"."COUNTRY_EN" "COUNTRY_EN", "Sales_territory_dimension"."FLAG_IMAGE" "FLAG_IMAGE31", "Sales_territory_dimension"."SALES_TERRITORY_EN"
"SALES_TERRITORY_EN"
from "GOSLDW"."SALES_TERRITORY_DIMENSION" "Sales_territory_dimension"), "Gender_lookup13" as (
select "Gender_lookup"."GENDER_CODE" "GENDER_CODE", min("Gender_lookup"."GENDER") "GENDER"
from "GOSLDW"."GENDER_LOOKUP" "Gender_lookup"
where "Gender_lookup"."LANGUAGE" = 'EN'
group by "Gender_lookup"."GENDER_CODE"), "Retailer__model_" as (
select "Retailer_dimension11"."RETAILER_SITE_KEY" "Retailer_site_key", "Sales_territory_dimension12"."SALES_TERRITORY_KEY" "Sales_territory_key",
"Sales_territory_dimension12"."SALES_TERRITORY_EN" "Sales_territory"
from "GOSLDW"."RETAILER_DIMENSION" "Retailer_dimension11", "Sales_territory_dimension12", "Gender_lookup13"
where "Retailer_dimension11"."GENDER_CODE" = "Gender_lookup13"."GENDER_CODE" and "Retailer_dimension11"."COUNTRY_KEY" = "Sales_territory_dimension12"."COUNTRY_KEY")
select "Order_method_dimension14"."ORDER_METHOD_KEY" "Order_method0key", "Order_method_dimension14"."Order_method" "Order_method1",
"Product_line15"."PRODUCT_LINE_CODE" "Product_linekey", "Product_line15"."Product_line" "Product_line0", "Retailer__model_" "Sales_territory_key" "Retailer_territorykey",
"Retailer__model_" "Sales_territory" "Sales_territory", cast("Time_dimension17"."CURRENT_YEAR" as char(4)) "Yearkey", cast("Time_dimension17"."QUARTER_KEY" as char(6)) "Quarterkey",
cast("Time_dimension17"."MONTH_KEY" as char(6)) "Monthkey", sum("Sales_fact18"."GROSS_PROFIT") "Gross_profit" from "Order_method_dimension14", "Product_line15", "Retailer__model_",
"GOSLDW"."TIME_DIMENSION" "Time_dimension17", "GOSLDW"."SALES_FACT" "Sales_fact18", "GOSLDW"."PRODUCT_TYPE" "Product_type19", "GOSLDW"."PRODUCT_DIMENSION"
"Product_dimension20"
where "Order_method_dimension14"."ORDER_METHOD_KEY" = "Sales_fact18"."ORDER_METHOD_KEY" and "Product_dimension20"."PRODUCT_KEY" = "Sales_fact18"."PRODUCT_KEY" and
"Product_type19"."PRODUCT_TYPE_CODE" = "Product_dimension20"."PRODUCT_TYPE_CODE" and "Product_line15"."PRODUCT_LINE_CODE" = "Product_type19"."PRODUCT_LINE_CODE"
and "Time_dimension17"."DAY_KEY" = "Sales_fact18"."ORDER_DAY_KEY" and "Retailer__model_" "Retailer_site_key" = "Sales_fact18"."RETAILER_SITE_KEY"
group by "Order_method_dimension14"."ORDER_METHOD_KEY", "Order_method_dimension14"."Order_method", "Product_line15"."PRODUCT_LINE_CODE", "Product_line15"."Product_line",
"Retailer__model_" "Sales_territory_key", "Retailer__model_" "Sales_territory", cast("Time_dimension17"."CURRENT_YEAR" as char(4)), cast("Time_dimension17"."QUARTER_KEY" as char(6)),
cast("Time_dimension17"."MONTH_KEY" as char(6))
```


BLU Acceleration Beats Competitor on the SAME Hardware in Single User Tests



DB2 10.5 with BLU Acceleration

40 Intel Westmere cores
512GB RAM
4x500GB HDD
RHEL 6.4, DB2 v10.5



3YR TCA
\$186,000

220GB Serial Execution workload

10.78 sec
Intermediate Reports

523 sec
Complex Reports


- ✓ **11.5x Faster**
for Intermediate reports
- ✓ **2.9x Faster**
for Complex reports
- ✓ **1/12 the Cost**

BI Day Deep Analytics Test Measures Throughput for Concurrent Mix of Data Intensive Queries

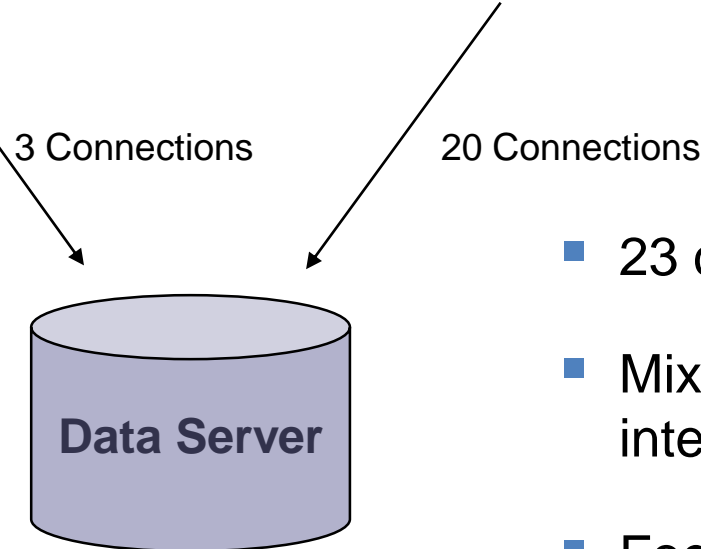
3 Users doing **complex** reports



20 Users doing **intermediate** Reports



Each report executes one or more queries



- 23 concurrent users
- Mix of complex and intermediate reports
- Focus on throughput

Note: Distribution of complex, intermediate, and simple workloads based on Forrester Research, Profiling the Analytic End User for Business Intelligence

BLU Acceleration on POWER8 S824 Beats Pre-Integrated Database Competitor V4 in **Deep Analytics** Workloads

IBM POWER8 S824



24 core Power8
256GB RAM
AIX 7.1 64bit

Flash System 840



3YR TCA
\$616,276

1TB Deep Analytics workload











27.8 RPH

Mean Reports/Hour

- ✓ **1.4x More** reports/hour
- ✓ **6.2x Better** Price/Performance

Actionable Compression Minimizes Storage, Reduces I/O and Maximizes Performance

- Approximate Huffman encoding yields great compression
 - The more frequent a value is repeated, the fewer bits needed to store

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

- Encoding is **order preserving** so data is “**actionable**” in its compressed form

Only BLU Acceleration Enables Query Processing Directly on Compressed Data

- **Actionable Compression:** Encoded values do not need to be decompressed during query evaluation
 - Predicates, group by, aggregations, joins, etc work directly on encoded data

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

LAST_NAME	Encoding
-----------	----------

Brown	
Johnson	
Johnson	
Johnson	
Johnson	
Brown	
Johnson	
Gilligan	
Wong	
Johnson	



Oracle Exadata Hybrid Columnar Compression Sacrifices Performance

- A compression technique, **not a columnar database**
 - A set of row oriented pages are rearranged, grouping columns together, then compressed into a ‘compression unit’
 - All columns for any given row are stored in a “compression unit”, which is the smallest unit of I/O
 - No ability to read just one column value for many rows
- **Not actionable** – data is decompressed on storage and sent to database nodes for query processing
 - During Exadata “smart scan”, decompression process consumes CPU resources in the storage servers
 - Rows are reconstructed out of the compression units
 - Uncompressed rows are returned to database servers

BLU Next Generation In-Memory Processing: Ultra Efficient Memory Usage for Analytics

- Your data is growing quickly, even with great compression your data will still be larger than available RAM
- BLU dynamic in-memory processing means performance is still great even when data is many times larger than RAM
 - **Data does not need to fit into RAM**
- New algorithms for prefetching and caching compressed column data
- Only column data relevant to queries occupies RAM, other data stays on disk until it is needed
 - Analytical queries typically access only a subset of columns
 - A small amount of RAM can cache a very large percentage of relevant data
 - Column data is moved in and out of cache in pages

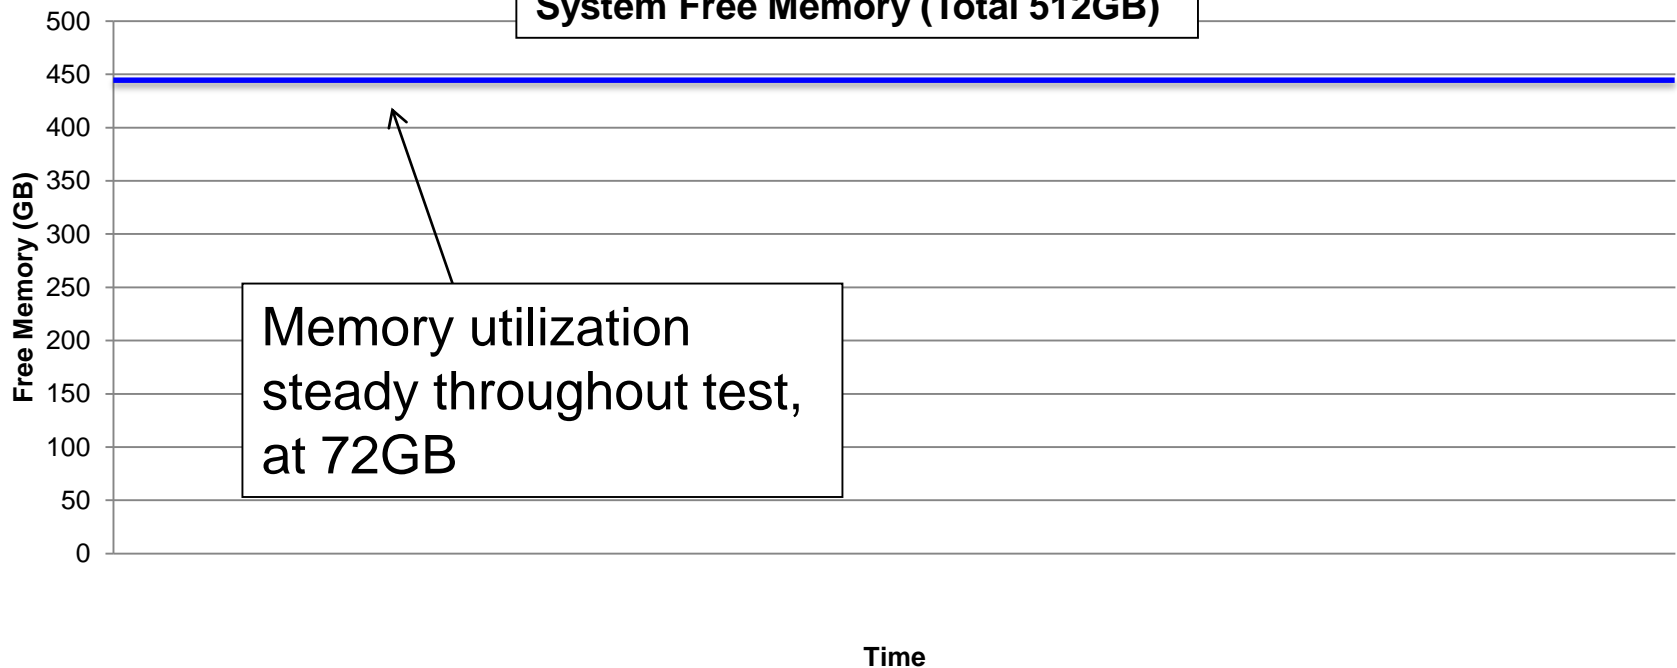
BLU Acceleration Utilizes Memory Efficiently and Predictably

Graph of Free Memory during a Single User Test

- 96GB of compressed data (550GB uncompressed)
- System “primed” prior to test to load data into memory

DB2 10.5 x3850 (512GB total)

System Free Memory (Total 512GB)



Is SAP HANA Enterprise Ready?



Release Date	Release Number	Days from last release
8-Jan-13	47	
14-Jan-13	48	6
23-Jan-13	49	9
8-Feb-13	50	16
19-Feb-13	51	11
11-Mar-13	52	20
5-Apr-13	53	25
25-Apr-13	54	20
29-Apr-13	55	4
15-May-13	56	16
11-Jun-13	57	27
13-Jun-13	58	2
2-Jul-13	60	19
21-Jul-13	61	19
5-Aug-13	62	15
22-Aug-13	63	17
4-Sep-13	64	13
13-Sep-13	65	9
14-Sep-13	66	1
27-Sep-13	67	13
17-Oct-13	68	20
12-Nov-13	69	26
4-Dec-13	70	22
16-Feb-14	71	74
23-Feb-14	72	7
17-Mar-14	73	22
15-Apr-14	74	29
28-May-14	80	43
3-Jul-14	81	36
22-Aug-14	82	50

- SAP HANA's release frequency is ***PAINFUL***

- AVERAGE RATE = ***every 20 days***

 28 hours

DB2 with BLU Delivers Clear Advantages for SAP Environments



Accelerated Performance

- Over 2x better, 24 core performance for SAP SD benchmark than nearest competitive result¹



Faster Time to Value

- Simple load & go in-memory support
- No change to applications needed

Business Proven



Lower Risk

- Accelerates SAP BW without application upgrade
- 10+ year track record of SAP optimization



Transparent Scalability

- Active data can be larger than available memory
- Automated workload management with shared-everything resource pools

Lower Cost



Lower Acquisition Costs

- Leverage existing investments
- No need to buy additional HW



Lower Operating Costs

- Option for lower cost, near line storage
- Mature technology not requiring frequent patching

(1.0) IBM Power System S824 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 4 processors / 24 cores / 96 threads, POWER8; 3.52GHz, 512 GB memory, 21,212 SD benchmark users, running AIX® 7.1 and DB2® 10.5, dialog response: 0.98 seconds, line items/hour: 2,317,330, dialog steps/hour: 6.952,000 SAPS: 115,870 database response time (dialog/update): 0.011 sec / 0.019sec, CPU utilization: 99%, Certification #: * Results valid as of 3/24/14. * Certification # not available at press time. Source: <http://www.sap.com/benchmark>. (1.1) Fujitsu RX300 S8 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 2 processors / 24 cores / 48 threads. Intel Xeon E5-2697 processor 2.70 GHz, 256 GB memory, 10,240 SD benchmark users, running Windows Server 2012 SE and SQL Server 2012, Certification #: 2013024

Oracle In-Memory – Late, Costly, Complex

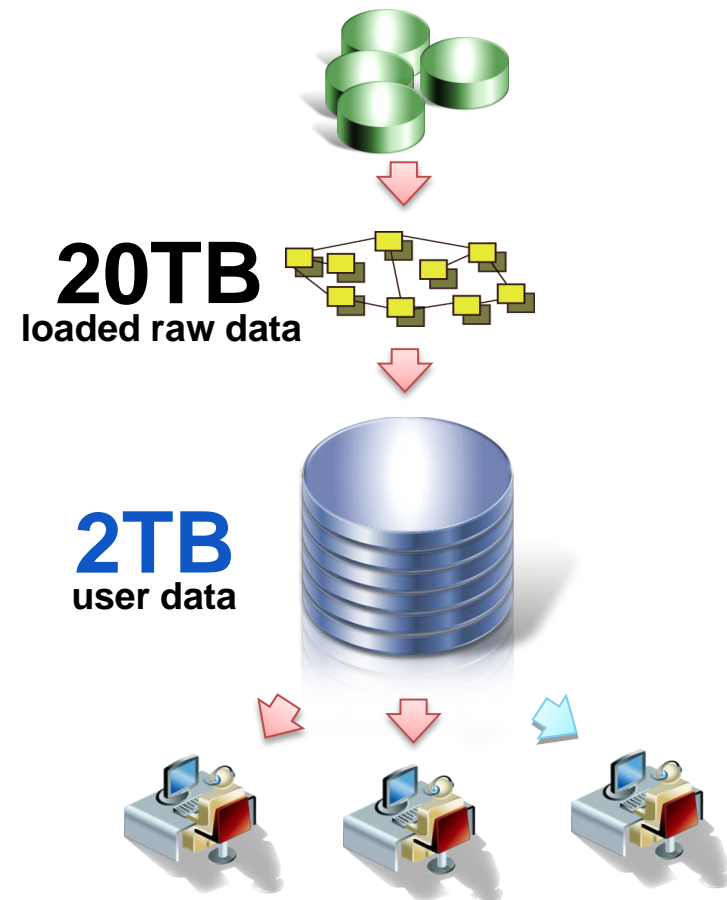
	BLU Acceleration	Oracle In-Memory Cache
Available Now?	Yes. Delivering value to clients since 2Q 2013	Recently released for Linux and Solaris in July 2014
Column Format Persistence	Yes. Fully persisted on disk and cached in memory	No. Only cached in memory. Must be fully recreated at a database restart.
Can Data be Larger than Memory?	Yes. Truly Load and Go	No. Yet another option. DBAs must decide when/how to use and must manage memory carefully. Not simple
Superior Storage Savings	Yes. 10x or higher compression. Compressed on disk and in memory	No. Store the data twice in memory. Claim 2x-4x compression in memory while requiring full table to be stored on disk with indexes
Actionable Compression	Yes. The broadest range of operations can act on compressed data today	No. In fact, Oracle documentation talks about the performance impact for choosing higher levels of compression in-memory
Added Cost	Included in both AESE and AWSE	Yes. Yet another add on, only available on top of Enterprise Edition

Terabyte Pricing for DB2 with BLU Acceleration

Much more power for much lower cost

Terabyte Pricing:

- Pay for only the data you analyze
- Leverage compression to reduce the amount of licensed data (e.g. 20TB compressing to only **2TB**)
- Independent of hardware
- Low, predictable licensing cost
- Add more cores and memory without affecting the licensing cost



BLU Acceleration is Available in the Cloud

BLUforCloud.com

DB2 BLU for Cloud

Self-service Data Warehousing & BI in the Cloud

The same benefits of DB2 with BLU Acceleration, plus

- Complete analytics solution, deployed in under an hour
- Powerful database & analytic capabilities at a fraction of the cost
- No infrastructure investment
- Cognos Business Intelligence included
- Expert-built schemas for business insight

Available on:



SOFTLAYER[®]
an IBM Company

For use by:

Business	Data Scientists Business Analysts Line of Business Users
IT	DBAs Developers