



DB2 with BLU on Power Analytics at the Speed of Thought

Gürşad Küçük

CTP Leader of DB2 for SAP
IBM Software Group Europe

ISICC - IBM SAP International Competence Center
69190 Walldorf, Germany





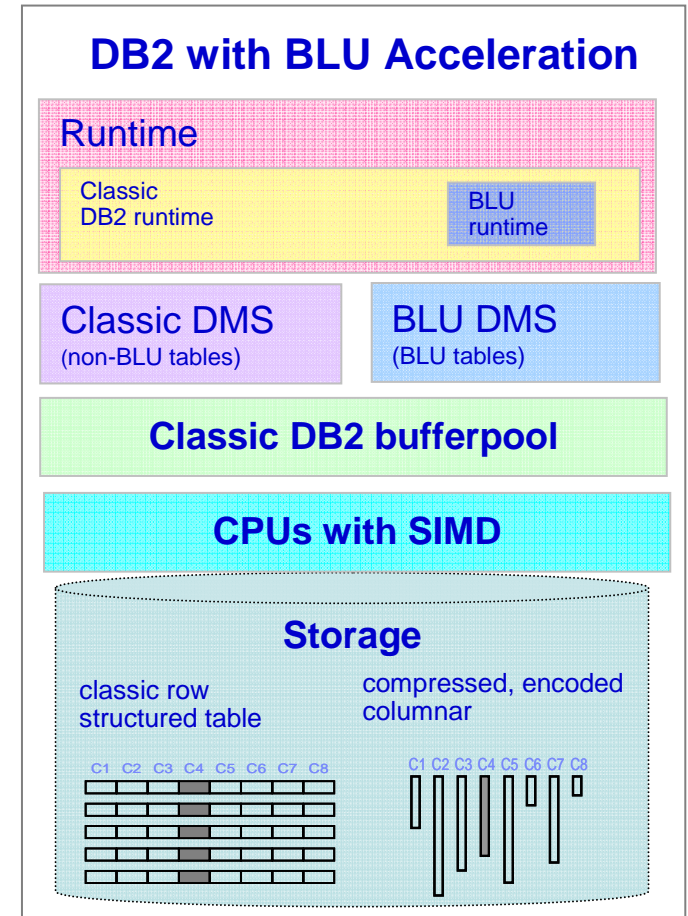
Agenda

- **DB2 BLU Architecture**
- **DB2 BLU Advantages for POWER Servers**
- **DB2 BLU POC Results**



BLU is Built Into DB2 Engine

- **Blink accelerator technology developed by IBM Almaden Research since 2007**
 - Predictably fast, interactive and ad hoc querying
 - User need not to change SQL or application code
- **DB2 BLU (BLink Ultra) technology for accelerating queries**
 - CPU-optimized unique runtime handling
 - Unique encoding for speed and compression
 - Unique memory management
 - Columnar storage, vector processing (e.g. SIMD)
 - Built directly into the DB2 kernel
- **Revolution and evolution**
 - BLU tables coexists with traditional row tables - in same schema, storage and memory
 - Query any combination of row or BLU tables
 - Easy conversion of tables to BLU tables
 - Change everything at one, or change incrementally
- **DB2 capacity measured in Peta and Zeta Bytes**
 - 1 DB2 partition: 2048 PB (uncompressed ~8000-20000 PB)
 - 1000 DB2 partitions: 2'048'000 PB (uncompressed ~8-20 Zeta Bytes)






DB2 BLU: Key Features for Performance Improvement

IBM Research & Development Lab Innovations

• Dynamic In-Memory

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



• Actionable Compression

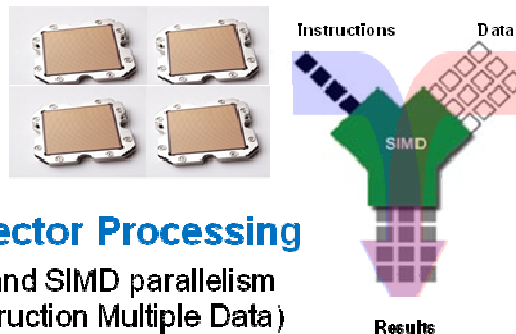
Patented compression technique that preserves order so that the data can be used without decompressing



Encoded


• Parallel Vector Processing

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



• Data Skipping

Skips unnecessary processing of irrelevant data



Super Fast, Super Easy — Create, Load and Go!
No Indexes, No Aggregates, No Tuning, No SQL changes, No schema changes



Estimated HW Infrastructure for Production – Year 1 and Year 5: assumption yearly 20% growth

- HW estimation considers database server and application server
- DB2 BLU implemented on 1 medium POWER server (2-tier)
- „In-Memory“ spread to many large Intel servers (3-tier)
- Source Oracle database 8 TB on BW 7.0 (non-unicode)

Year 1: Production	DB2 10.5 BLU	In-Memory	Difference in %
# cores	40	332	88.0%
RAM in GB	512	4166	87.7%
Storage in GB	2091	40960	94.9%
Database in GB	1608	16664	90.3%
Year 5: Prod. with 20% growth	DB2 10.5 BLU	In-Memory	Difference in %
# cores	40	826	95.2%
RAM in GB	512	10366	95.1%
Storage in GB	5203	101922	94.9%
Database in GB	4002	41465	90.3%

**Huge savings
through DB2
technology**

**→ DB2 BLU
needs approx.
90% less HW**

Remark: DB2 BLU is not limited to the number of cores and RAM capacity
→ database storage size can grow with same server capacity



DB2 BLU Performance – Less RAM than Data ?


SAP BW scenario

- Fact table with 438 Million records (~200GB)
- Size (table + indexes)
 - Row based (adaptive compr.): 33 GB
 - **BLU compressed: 12 GB**
- DB2 parallel query degree switched ON

DB2 10.5 Test environment

- 8 cores,
- 32 GB RAM
- SAP BW 7.30

Test 1: 15 GB Bufferpool
Test 2: 7.5 GB Bufferpool

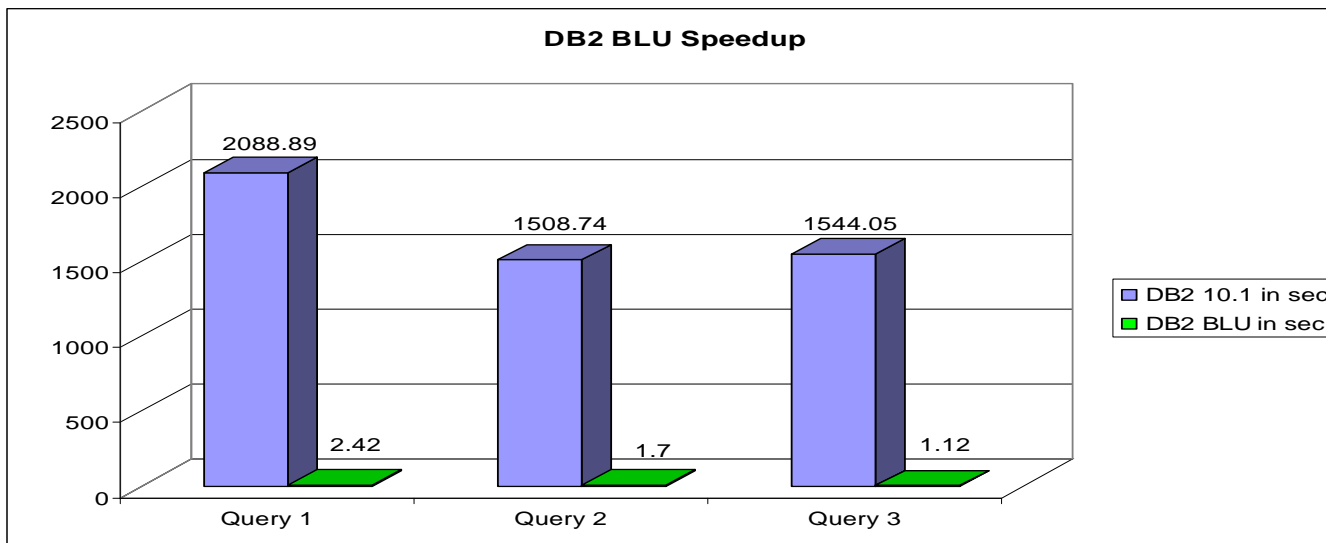


DB2 BLU runs exceptionally well even with less RAM



DB2 BLU Performance for Warehouse Queries

	DB2 10.1 in sec	DB2 BLU in sec	Speedup over DB2 10.1
Query 1	2088.89	2.42	863
Query 2	1508.74	1.7	887
Query 3	1544.05	1.12	1379



- DB2 BLU on same HW like DB2 10.1
- DB2 BLU uses advantages and architectures coming from DB2 10.1
- DB2 BLU optimizes the existing techniques in DB2 10.1



DB2 10.5 BLU: Evaluation for Operations

<i>Criteria compared with uncompressed source system</i>	DB2 10.5 / BLU
OS support	AIX, Linux (Win planned)
SAP release	7.0 and higher
Virtualisation (incl. production)	Yes
2-Tier support	Yes
3-Tier support	Yes
NLS (nearline storage) support as underlying database	Yes
SAP support ends	31.12.2022
Number of patches per year ****	~2
HW / IT support	Appliances and tailored server
Percentage of source storage *	~10 to 30%
Percentage of source RAM **	~100 to 200%

* DB2 compression reduces storage capacity by 70-95%, depending on share of column-store objects

** DB2 9.7/10.1 requires usually 2-5% of the database size, with DB2 10.5 4-10%

**** based on experiences with DB2 V8, V9, DB2 10.1, DB2 10.5

Status: March 2014



Agenda

- DB2 BLU Architecture
- **DB2 BLU Advantages for POWER Servers**
- DB2 BLU POC Results



In-Memory: RAM is fast, Cache is faster, Register is fastest

Type of Memory	Size	Latency
L1 cache	64 KB	~4 cycles [2 ns]
L2 cache	256 KB	~10 cycles [5 ns]
L3 cache (shared)	8 MB	35-40+ cycles [20 ns]
Main memory	GBs up to terabytes	100-400 cycles
Solid state memory	GBs up to terabytes	5,000 cycles
Disk	Up to petabytes	1,000,000 cycles

cache, register
RAM,
storage

25-100

10 k

Today (without in-memory technology): Approx. 100% of the application & database data **in RAM** !

Challenge: maximize usage of cache & register by using cores more efficient

Options:

- Database partitioning
- In-memory and column-store technology



DB2 BLU's View on POWER and Intel

	POWER 7+	x86 (i7: E78870)
Cores per Chip	8 cores	10 cores
Processor Speed	4.4 GHz	2.4 GHz
Cache		
L1 Cache	32+32 kB/core X8 (512kB)	480 kB
L2 Cache	256 KB per C1 Core X8 (2 MB)	2 MB
L3 Cache	10MB per C1 Core X8 (80 MB)	30 MB (shared)
L4 Cache		
Threads per Core	4 Threads	2 Threads

POWER 7+: Fewer but more efficient cores with 4 threads per core as compared with 2 for Intel

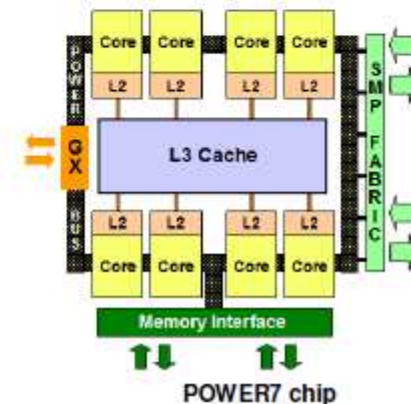

- much more cache
- faster processors

POWER 8: Will provide access to a lot (tens to hundreds of terabytes) of Flash memory – enabling to build extremely large, extremely fast DB2 BLU databases



DB2 Exploits Unique Features of POWER

- POWER6 on core Decimal Floating Point Unit
- DB2 is the only DBMS vendor to natively support `DECFLOAT` data type
 - Performance advantage for retail and finance
 - 40% performance gain in SAP BW
 - Have seen **up to 6x faster performance**
- `DECFLOAT` is internally used by DB2 Oracle compatibility mode for `NUMBER` data type.
 - DB2 Oracle compatibility is used for example for SAP Identity Management product (SAP IDM)
- POWER7 chip has on-chip L2 / L3 cache with eDRAM L3 Cache
 - DB2 is cache aware
 - Optimizes power of the core
- POWER7 features used by BLU
 - Decimal Floating Point Unit (DFU) utilization for DECIMAL data type conversions
 - VSX vector unit
 - POWER7 instruction set fully exploited by DB2 10.5 BLU





Green IT: DB2 on POWER vs In-Memory on Intel

Customer runs DB2 on POWER

- 180 systems, 48 production
- 26 HA (LPM*) + 26 DR (PowerHA)
- 2 data centers

→ 4 POWER servers



Possible In-Memory on Intel **

- 180 systems, **48 production**
- **26 HA + 26 DR clusters**
- 2 **BIGGER** or more data centers
- 48 Intel servers for production
- 52 Intel servers for HA+DR clusters
- up to 48 Intel servers for test/QA
- up to 48 Intel servers for dev
- up to 36 Intel servers for rest

→ 101-232 Intel servers

4 POWER servers versus 101+ Intel servers !!!

* LPM - AIX live partition mobility

** SAP note 1788665, 1681092: No virtualization and no multiple In-Memory databases on a production appliance




Agenda

- DB2 BLU Architecture
- DB2 BLU Advantages for POWER Servers
- **DB2 BLU POC Results**



DB2 BLU PoC results

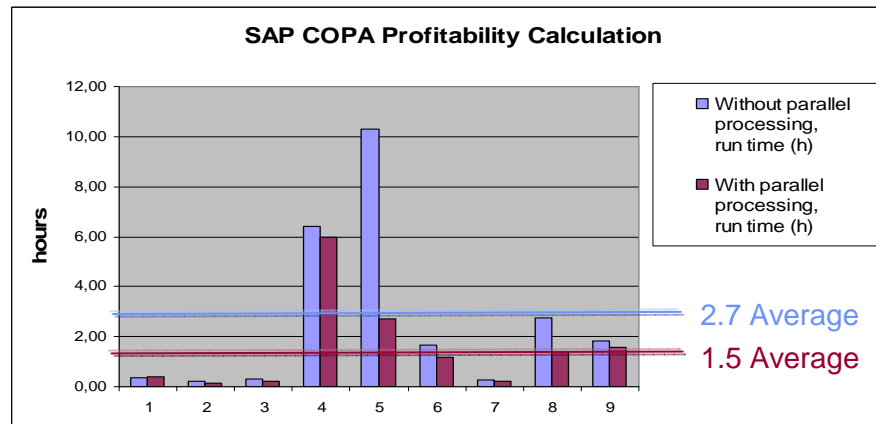
- DB2 BLU used between 50%-84,6% less cores
- DB2 BLU used in average 75% less cores and was in average 24% faster
- **DB2 BLU delivers with less server capacity similar or better performance**





DB2 BLU: SAP ERP / CO-PA Performance

- **DB2: SAP CO-PA performance improvement (First customer results with DB2 9.7, without BLU)**
 - up to factor 3.8x faster and in average 1.8x faster on SAME hardware and SAME database
 - with DB2 10.5 BLU further performance possible





DB2 10.5 Advantages for SAP Applications over other In-memory Solution

Strategic Partnership with SAP and DB2

- SAP involved in the DB2 development
- **SAP runs over 1.000 systems on DB2**
- DB2-SAP roadmap/support until 2022
- **SAP 4.6 and higher supported with DB2 9.7+**
- DB2 follows SAP's „7+2“ years support strategy
- **Adjusted development with SAP**
- Fast support often in 4-8 weeks
- **Reduced requirement for database upgrades/patches**
- Less tuning required (DBACOCKPIT, SAP knob)

Reducing the Overall Costs (TCO)

- Reduced DB storage 80%-95%
- Reduced RAM for database 50-90%
- **Reduced cores for database 50-90%**
- Reduced DBA administration > 50%
- Reduced energy costs (virtualisation) up to 80%
- Reduced patch cycles up to 25 x
- Reduced upgrade cycles EOS Dec-2022
- Reuse of existing HW non-disruptive
- BLU, HADR, pureScale, NLS incl in AESE

Innovative Database Technology

- In-memory database optimization at core and cache level
- **Row-store and column-store can co-exist on same server and same storage**
- One database engine for OLTP/OLAP/NLS
- Leading DB2 performance for various official benchmarks (e.g. SAP SD and TPC benchmark)
- **Integrated high-availability and disaster-recovery**
- RTO & RPO → 0
- **VMWare, XEN and Hyper-V fully supported**

Reference Customers

- **Approx. 20.000 SAP systems on DB2**
- **Eli Lilly:** Redhat-Linux, VMware and EMC storage. Better performance, over 62% storage saving (during unicode conversion)
- **AUDI:** AIX and IBM DS8K storage. Better performance, up to 74% storage saving
- **Coca-Cola:** DB2 10.5 BLU on AIX. Tables become 10-25x smaller and analytical query ran 4-15x faster
- **BNSF Railway:** DB2 10.5 BLU. On average analytic queries running 74x faster
- **DB2 BLU:** <http://www.youtube.com/watch?v=5T6f74gYu1Y>

IBM®

Teşekkürler

