

Why Power is a platform of choice for Databases and Data Warehouses



Mladen Jovanovski

Technical Sales Specialist

Information Management Software

IBM Software Group SEE

mladen.jovanovski@rs.ibm.com

Agenda

- **DB2 advantages on Power**
- **DB2 Data Warehouse on Power**
- **DB2 High Availability on Power**
- **DB2 BLU Acceleration**
- **PureData Systems**

DB2 advantages on Power



Power Systems and DB2 – Tight Integration

- Performance of POWER7+ and DB2
 - More cores and threads – 32 chips, 8 cores/chip, 4 threads/core
 - Exploited by DB2 better than any other DBMS
 - Full SSD support in Power Systems
 - DB2 can use SSD for both permanent objects (tables/indexes) as well as temporary objects (not supported in Exadata)

- Consolidation of DB2 on POWER7+
 - PowerVM virtualization second to none
 - Active Memory Sharing exploited by DB2 self tuning memory manager
 - When peak demand hits, DB2 can immediately leverage additional resources
 - Workload management integrated between AIX and DB2
 - DB2 able to meet customers SLAs more easily

- Reliability of Power Systems and DB2
 - Power 3x – 4x more reliable than Linux on x86
 - 99.997% availability with Power and AIX
 - DB2 tightly integrated with PowerHA and other HA features of AIX

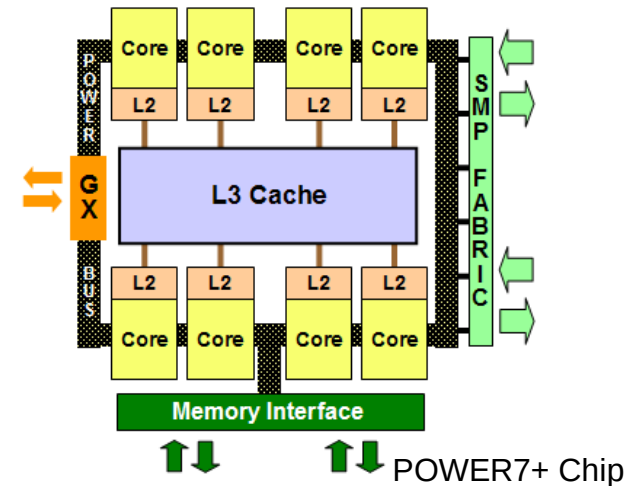
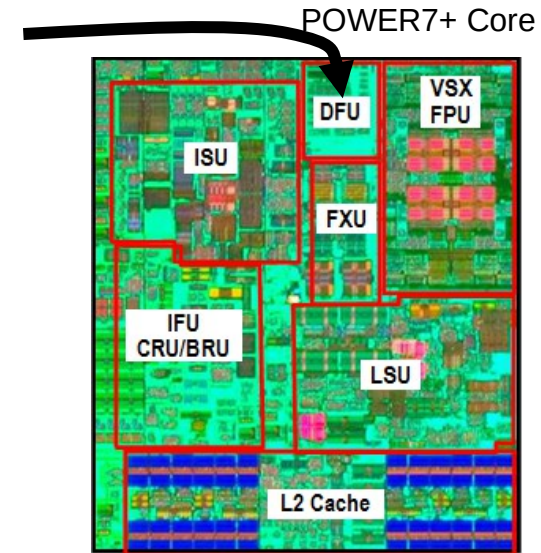
DB2 Architecture Maps to POWER7+ Capabilities

- **POWER7+ – massive number of threads per server**
 - Requires sophisticated software to **exploit**
 - DB2 threaded engine built to scale on large multi core servers
 - Requires sophisticated **virtualization** to consolidate
 - DB2 autonomics “play nice” and “react quickly” in virtualized, dynamic environments
 - Requires advanced **workload management** to meet SLAs
 - DB2 and AIX tightly integrated WLM to deliver the resources where they are needed most
 - Requires **advanced diagnostics** to help lower administration costs for customers with massive levels of concurrency
 - The blue stack helps resolve problems faster with integrated diagnostics
 - Requires integrated **high availability**
 - If any part of the solution fails, DB2 and PowerHA respond more rapidly to provide business continuity

Unique Features of POWER7+ Core Only DB2 Exploits

- **POWER7+ 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

- **POWER7+ chip has on-chip L2/L3 cache with eDRAM L3 Cache**
 - 10 MB per core / 80 MB per chipset
 - DB2 is cache aware
 - Optimizes power of the core



Technology Driving Performance Advances

- **Process Exploitation**
 - Deep exploitation of Simultaneous Multi Threading (SMT)
 - Fully threaded DB2 engine
 - NUMAtization of DB2 resources to align with system architecture

- **Memory Exploitation**
 - Autonomic exploitation of POWER features such as larger page sizes
 - Support for AIX multi page support that includes 64KB, 16MB and 16GB AIX page sizes
 - Co-operative Caching

- **Storage Exploitation**
 - Exploits Asynchronous I/O and Scatter / Gather I/O, as well as AIX CIO and DIO interfaces
 - End-to-End I/O Priorities
 - Atomic Logical Volumes

- **Enablement for POWER6 features (Decimal Floating Point, Storage Keys)**

- **Deep integration with AIX APIs**

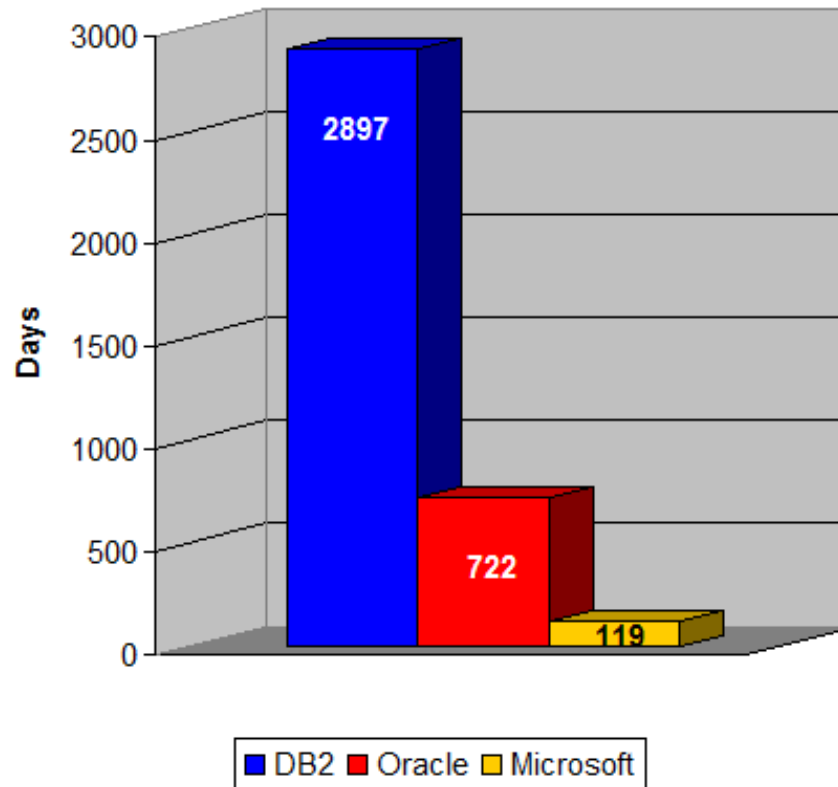
- **Exploits xIC capabilities for optimal performance using Profile Directed Feedback**

- **And many, many more ...**

Longevity in Transaction Processing Performance

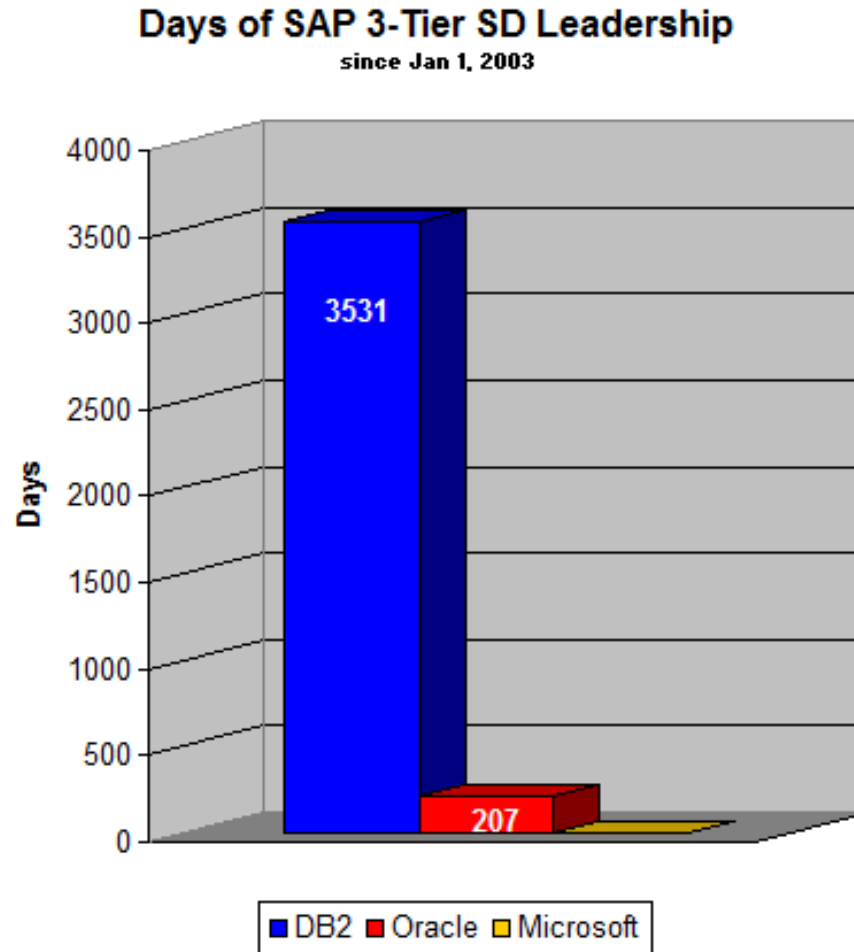
- Benchmarks are often a game of leapfrog
- DB2 has demonstrated sustained superiority over competitors

Single Server Days of TPC-C Leadership
since Jan 1, 2003



Longevity in SAP 3-Tier SD Performance

- DB2 has been unchallenged in SAP 3-Tier SD for more than 8 years



DB2 Data Warehouses on Power



What's Important for Data Warehousing?

- **Fast data access means fast performance**
 - DB2 with IBM Power lead all vendors in TPC-H days of leadership

- **Intelligent optimizer**
 - At 1TB and above any mistake is a big one
 - DB2 has the most mature optimizer in the business

- **Scalability architecture that delivers linear scalability**
 - The database partitioning feature provides proven near-linear scalability for queries and utility operations such as load

What's Important for Data Warehousing?

- **Ease of Growth**

- The IBM Balanced Warehouse provides a prescriptive scaling methodology for data warehousing

- **Balanced I/O throughput**

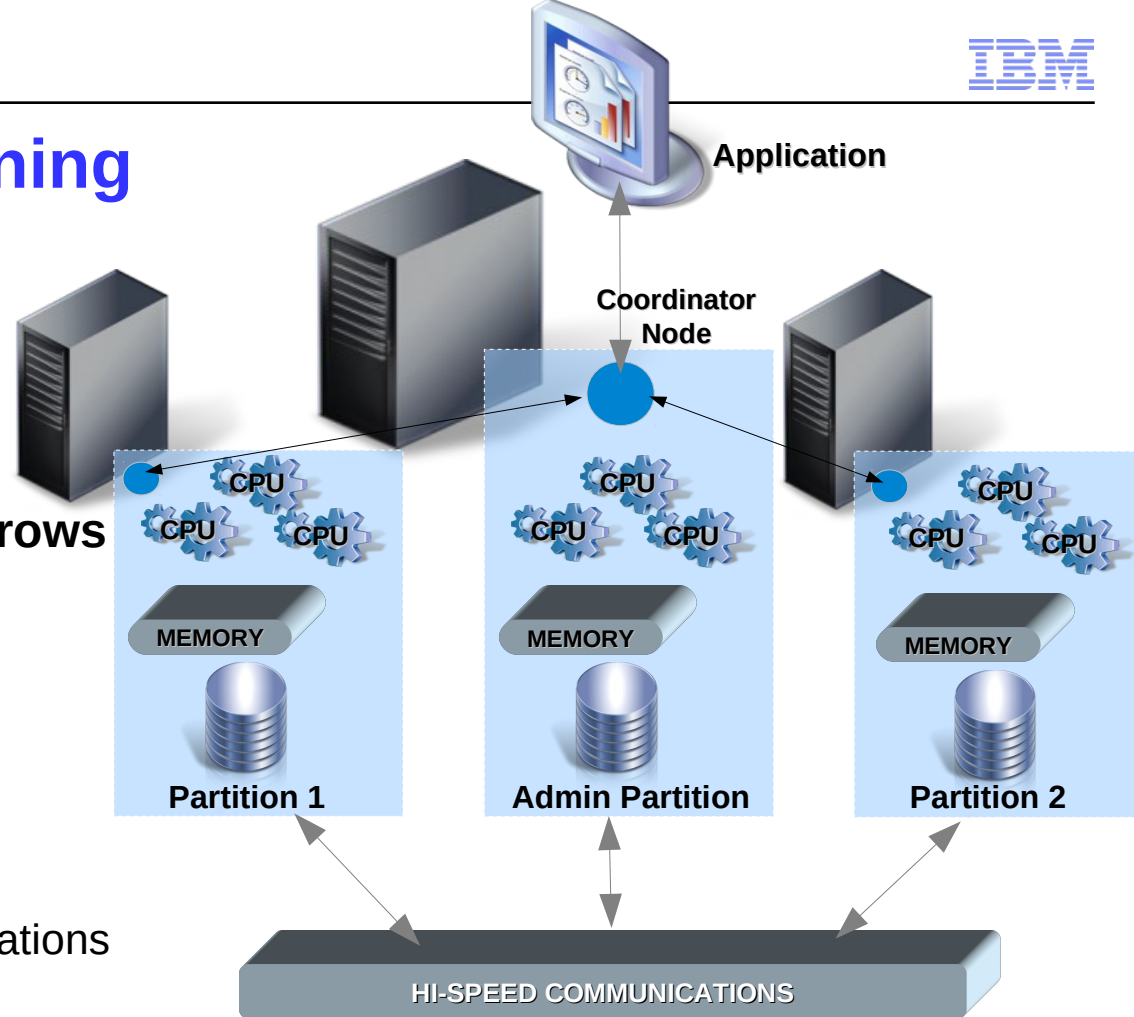
- IBM DB2 Warehouse removes the I/O bottleneck in a data warehouse
- DB2 has a number of cooperating technologies that minimize I/O and eliminate that bottleneck without the need to throw more processors (i.e. more cost) at the problem:
 - Multidimensional Cluster (MDC) tables
 - Storage Optimization Feature (Compression)
 - Database Partitioning Feature (DPF)

- **Workload Provisioning and Workload Management**

- DB2 has engine-level provisioning and management capabilities linked directly into AIX Workload Manager (WLM)

DB2 Database Partitioning Feature (DPF)


- Enables to evenly distribute rows across **database partitions (nodes)**.
- **Why Partition?**
 - Scale Out, Performance, ...
- **Benefits**
 - Transparent to users and applications
 - Parallelism (divide and rule)
 - Workload is divided among all nodes
 - Asynchronous I/O Parallel I/O
 - Dynamic throttling based on load
 - Near linear scalability
 - As the table grows, add more processing power in form of additional database partitions.



- DB2 core **Scale-Out** architecture based on **Parallelism** aka **Shared Nothing** architecture
 - Ability to spread all data across multiple server, each database partition has its own set of computing resources, including CPUs, memory, disk controllers and disks.

Longevity in Data Warehouse Performance

- First to publish 10TB TPC-H
- Only vendor to hold both 10TB TPC-H and leading TPC-C at the same time
- In the leapfrog game that is the world of Benchmarks, DB2 has longevity



DB2 High Availability on Power



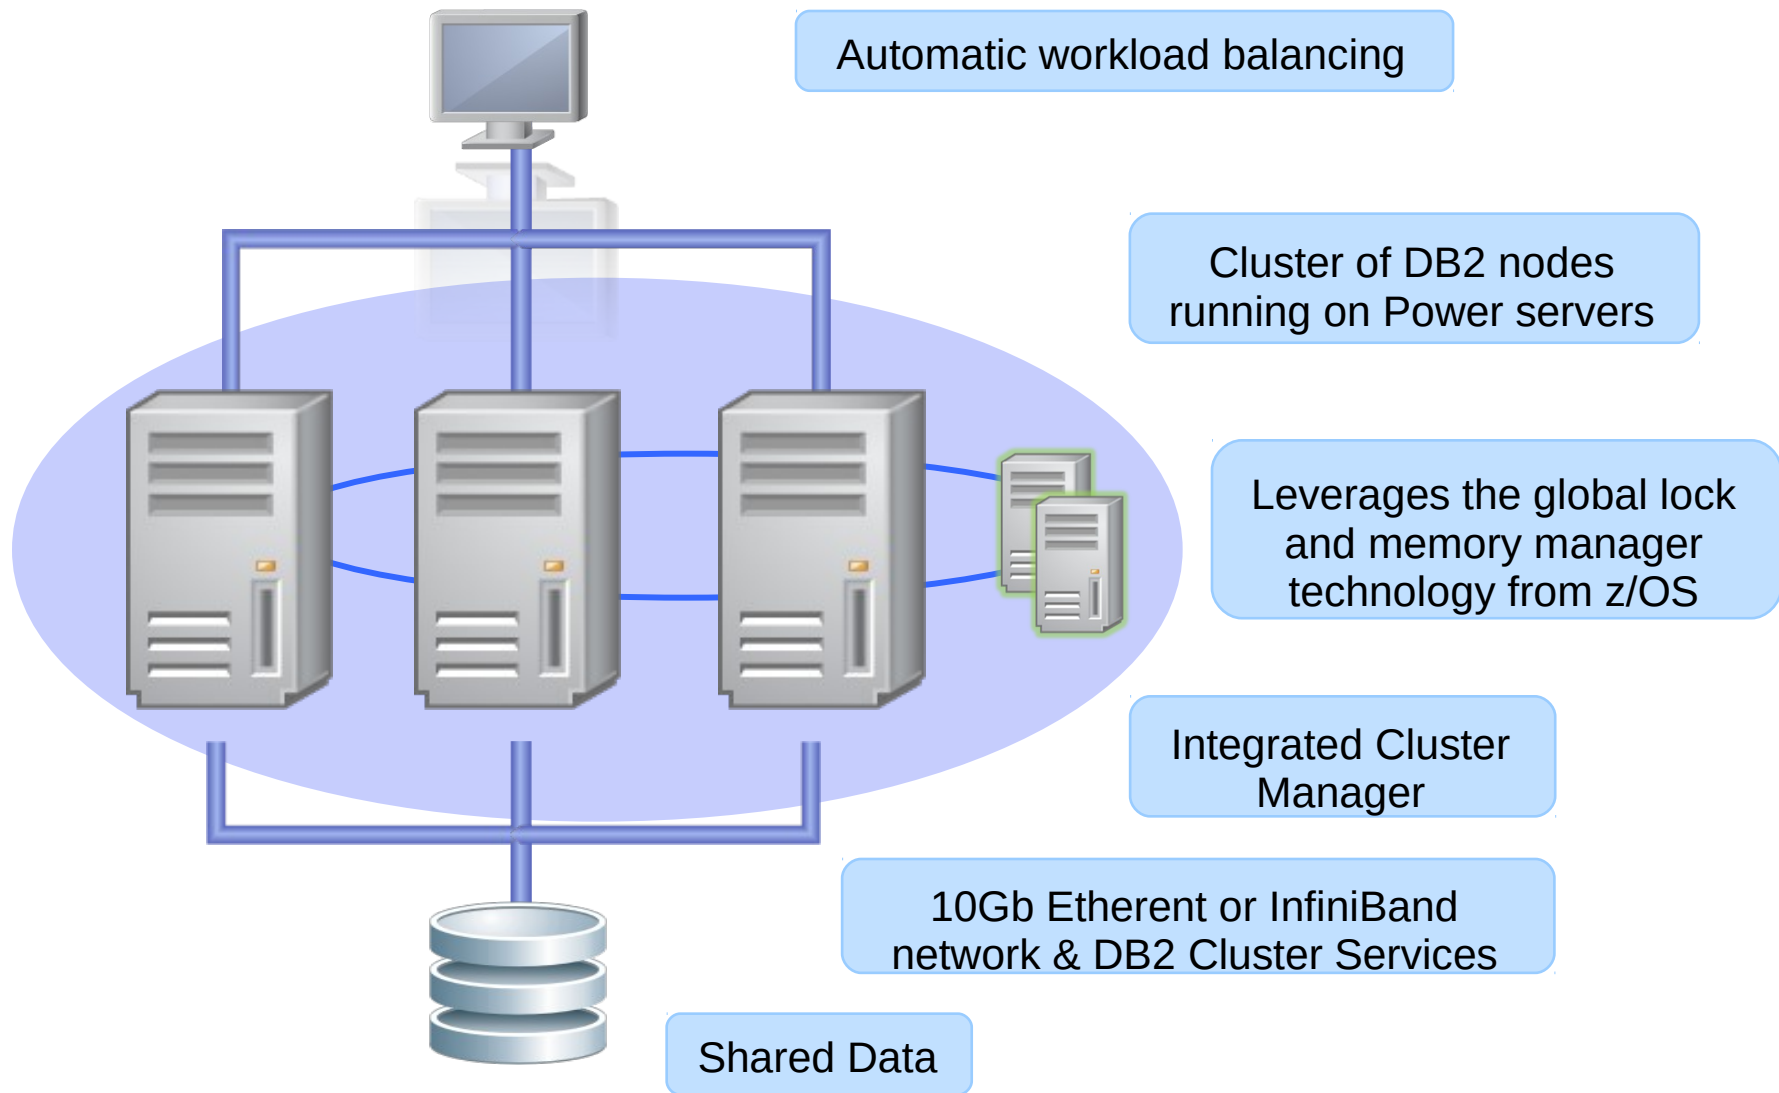
High Availability Options with DB2 on Power

- **Server failover**
 - Shared disk or remote disk mirroring

- **HADR**
 - HA and/or Disaster Recovery
 - Easy to set up and manage
 - Automatic failover with TSA integration
 - Fast failover

- **DB2 pureScale**
 - Unlimited Capacity
 - Buy only what you need, add capacity as your needs grow
 - Application Transparency
 - Avoid the risk and cost of application changes
 - Continuous Availability
 - Deliver uninterrupted access to your data with consistent performance

DB2 pureScale Architecture



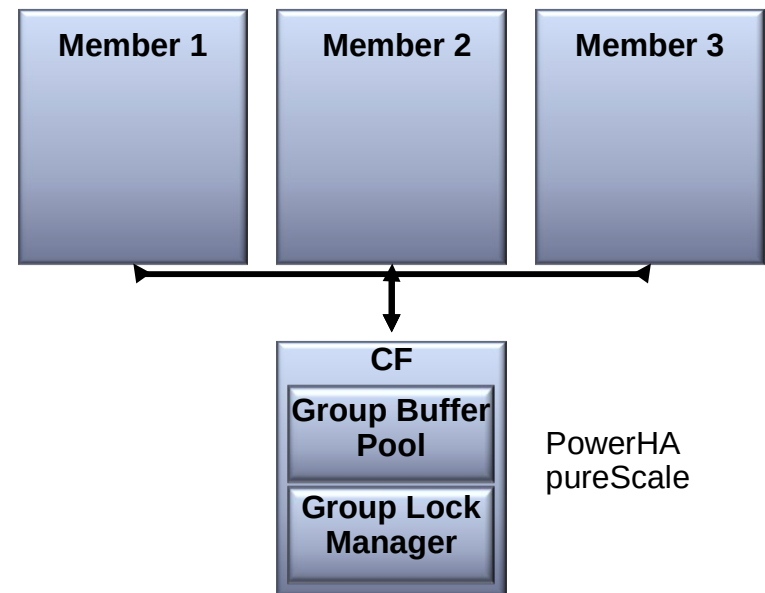
The Key to Scalability and High Availability

▪ Efficient Centralized Locking and Caching

- As the cluster grows, DB2 maintains one place to go for locking information and shared pages
- Optimized for very high speed access
 - DB2 pureScale uses Remote Direct Memory Access (RDMA) to communicate with the powerHA pureScale server
 - No IP socket calls, no interrupts, no context switching

▪ Results

- Near Linear Scalability to large numbers of servers
- Constant awareness of what each member is doing
 - If one member fails, no need to block I/O from other members
 - Recovery runs at memory speeds



DB2 BLU Acceleration

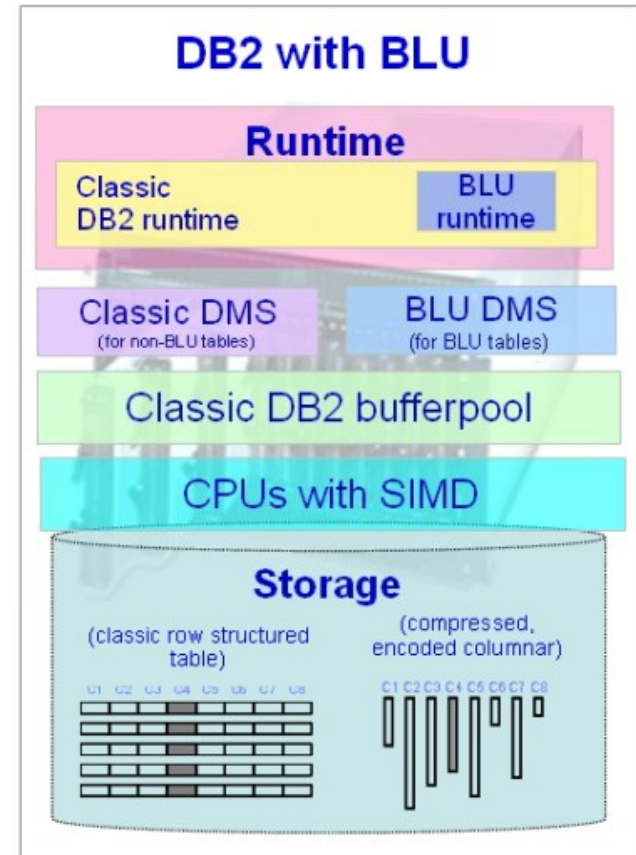


What is DB2 with BLU Acceleration?

- **Large order of magnitude benefits**
 - Performance
 - Storage savings
 - Time to value

- **New technology in DB2 for analytic queries**
 - CPU-optimized unique runtime handling
 - Unique encoding for speed and compression
 - Unique memory management
 - Columnar storage, vector processing
 - Built directly into the DB2 kernel

- **Revolution or 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, or change incrementally



Seamless Integration into DB2

- **Built seamlessly into DB2 – integration and coexistence**
 - Column-organized tables can coexist with existing, traditional, tables
 - Same schema, same storage, same memory
 - Integrated tooling support
 - Optim Query Workload Tuner recommends BLU Acceleration deployments

- **Same SQL, language interfaces, administration**
 - Column-organized tables or combinations of column-organized and row-organized tables can be accessed within the same SQL statement

- **Dramatic simplification – Just “Load and Go”**
 - Faster deployment
 - Fewer database objects required to achieve same outcome
 - Requires less ongoing management due to its optimized query processing and fewer database objects required
 - Simple migration
 - Conversion from traditional row table to BLU Acceleration is easy
 - DB2 Workload Manager identifies workloads to tune
 - Optim Query Workload Tuner recommends BLU Acceleration table transformations
 - Users only notice speed up; DBA’s only notice less work!
 - Management of single server solutions less expensive than clustered solutions

Simple to Deploy and Operate

DB2
WITH BLU
ACCELERATION



Super analytics
Super easy

- **Operations**
 - Simply **Load and Go**
 - Installation to business value in ~2 days
 - Ease of evaluation and performs as advertised

- **BI developers and DBAs – faster delivery**
 - No configuration or physical modeling
 - No indexes or tuning – out of the box performance
 - Data Architects/DBA focus on business value, not physical design

- **ETL developers**
 - No aggregate tables needed – simpler ETL logic
 - Faster load and transformation times

- **Business analysts**
 - Train of thought analysis – 5x to 100x faster
 - True ad-hoc queries – no tuning, no indexes
 - Ask complex queries against large datasets

Optimize the Entire Hardware Stack

In-Memory Optimized

- **Memory latency optimized for**
 - Scans
 - Joins
 - Aggregation
- **More useful data in memory**
 - Data stays compressed
 - Scan friendly caching
- **Less to put in memory**
 - Columnar access
 - Late materialization
 - Data skipping

CPU Optimized

- **CPU acceleration**
 - SIMD processing for
 - Scans
 - Joins
 - Grouping
 - Arithmetic
- **Keeping the CPUs busy**
 - Core friendly parallelism
- **Less CPU processing**
 - Operate on compressed data
 - Late materialization
 - Data skipping

I/O Optimized

- **Less to read**
 - Columnar I/O
 - Data skipping
 - Late materialization
- **Read less often**
 - Scan friendly caching
- **Efficient I/O**
 - Specialized columnar prefetching algorithm

PureData Systems




A new family of expert integrated systems

PureSystems

Systems with integrated expertise and built for cloud

Built-in Expertise
Capturing and automating what experts do – from the infrastructure patterns to the application patterns



Integration by Design
Deeply integrating and tuning hardware and software – in a ready-to-go workload optimized system

Simplified Experience
Making every part of the IT lifecycle easier - with integrated management of the entire system and a broad open ecosystem of optimized solutions

The IBM PureSystems Family

PureFlex

Infrastructure

Integrated and optimized infrastructure with flexibility

Runs your choice of applications and middleware




- Delivering IT infrastructure services

PureApplication

Application Platform

Integrated and optimized application platform

Built on IBM middleware to accelerate deployment of your choice of applications



- Delivering application platform services

PureData

Data Platform

Integrated and optimized data platform

Delivers high performance data services to transactional and analytics applications



- PureSystem with models optimized exclusively for data workloads

IBM PureData System: Optimized exclusively for data services


- **Optimized for data services**

- Transactional
- Analytics

- **Expert integrated**

- Data platform
- Infrastructure
- Unified platform management
- Built-in expertise

Data Platform



PureData

*Delivering
Data Services*

- **Data load ready in hours**

- **Data workload optimized**

- **Integrated management**

- **Single point of support**

- **Automated and integrated maintenance**

IBM PureData Systems built on Power

Meeting Big Data Challenges – Fast and Easy!



PureData

System for Transactions

For apps like E-commerce...

Database cluster services optimized for transactional throughput and scalability

PureData

System for Operational Analytics

For apps like Real-time Fraud Detection...

Operational data warehouse services optimized to balance high performance analytics and real-time operational throughput

IBM PureData System for Transactions


- **Speed**
 - Industry leading DB2 performance
- **Simplicity**
 - Data load ready in hours
 - Simplified system management
- **Scalability**
 - Highly reliable and scalable databases deployed in minutes
- **Smart**
 - Supports existing DB2 applications virtually unchanged
 - Supports existing Oracle Database applications with minimal or no change¹



Delivering highly scalable transactional data services

Built with DB2 pureScale Technology

- **Highly available and scalable database environment with application transparency**
- **Work spread across a cluster of nodes**
 - Acts as a single database system
- **If one node (member) fails, no disruption to work**
 - System *automatically* detects the failure
 - System *automatically* routes work to the other members
 - System *automatically* recovers and restarts the failed member
 - System *automatically* rebalances work across the members by default
- **Additional nodes can be added**
 - Start small and grow



IBM PureData System for Operational Analytics

Optimized exclusively for operational analytic data workloads

PureData



System for Operational Analytics

*Delivering data services for
operational analytics*



- **Speed**
 - Designed for 1000+ concurrent operational queries
 - Continuous ingest of operational data
 - MPP analytics (Massively Parallel Processing)
- **Simplicity**
 - Fast time-to-value
 - Automatic, policy-based data placement and workload management
 - Integrated management and support
- **Scalability**
 - Multiple sizes with data capacity up to Petabytes
- **Smart**
 - In-database analytics for leading applications
 - Supports DB2 applications unchanged and Oracle Database apps with minimal change
 - Clients have experienced cases of 10x storage space savings via Adaptive Compression

Built with DB2 Warehouse Technology (aka DPF)



- System Sizes – XS,S,M,L, XL+ Expansion Add-ons
- Single Pane of Glass for Management Console
- Raid Protected for All Storage
- 4x10G Ethernet per Data Module (active/active trunk)
- Consolidated Network up-links via reserve ports
- Roving High Availability support
- Fully redundant on all components by default

Summary

- **DB2 Performance on IBM Power**
 - Higher Performance = Lower Costs
- **DB2 High Availability with IBM Power**
 - More Reliable
- **DB2 and IBM Power Lowers Total Cost of Ownership**
- **Deep Technology Innovation Delivers More Value**
- **DB2 on POWER7+ delivers even more integration and more value to customers**