IBM Software Group

DB2 9 for LUW Update

RUG

Sandviken 23-24 jan 2007

mats.mohlin@se.ibm.com

ON DEMAND BUSINESS

© IBM Corporation

Agenda

- **Short DB2 V9 Update and V9 performance**
  - Berni Schiefer IBM DB2 Toronto lab
  - IOD conference in Demand Anaheim 2006
- **DB2 LUW trends and Direction**
  - **Application Development Strategy,**
    - Curt Cotner IBM DB2 SVL lab ?
    - IOD conference in Demand Anaheim 2006

IBM Information On Demand 2006

IBM

DB2 9 Performance

Berni Schiefer

IBM Toronto Lab

Presented by:
Mats Mohlin
mats.mohlin@se.ibm.com

TAKE BACK CONTROL

IBM INFORMATION ON DEMAND 2006
October 15 – 20, 2006
Anaheim Convention Center
Anaheim, California

Benchmarks – Triple Crown Winners

- DB2 is the performance leader
  - #1 TPC-C
    - 2.5x over Oracle
    - 3.2x over SQL Server
  - #1 SAP SD 3-tier
    - 1.7x over Oracle
    - 1.8x over SQL Server
  - #1 TPC-H @ 10TB
    - 1.7x over Oracle

http://www.ideasinternational.com/

http://www.tpc.org/

Version 8 to Version 9 Comparison

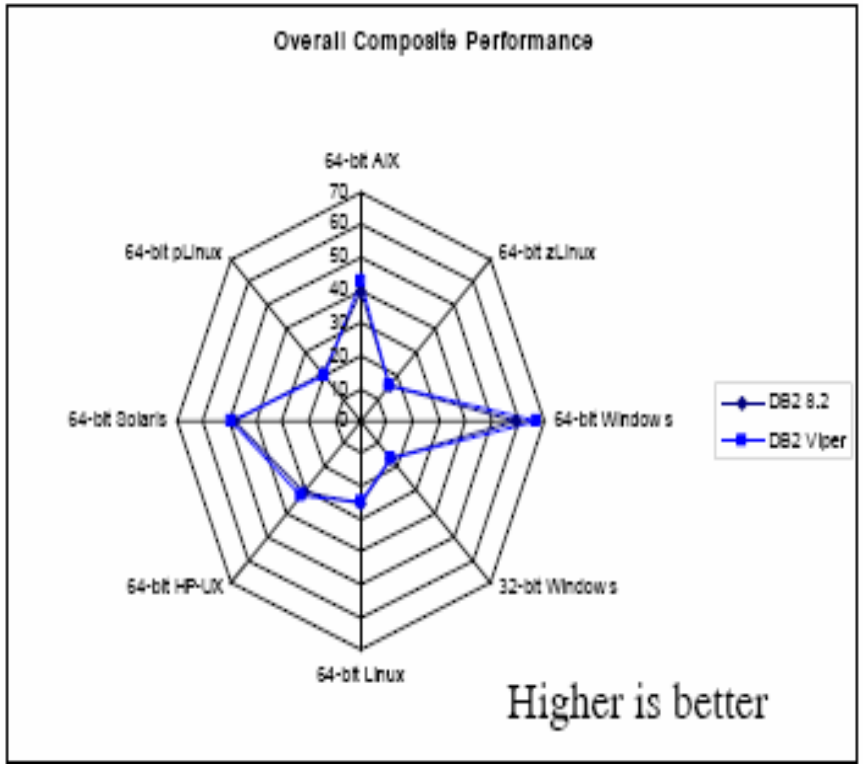

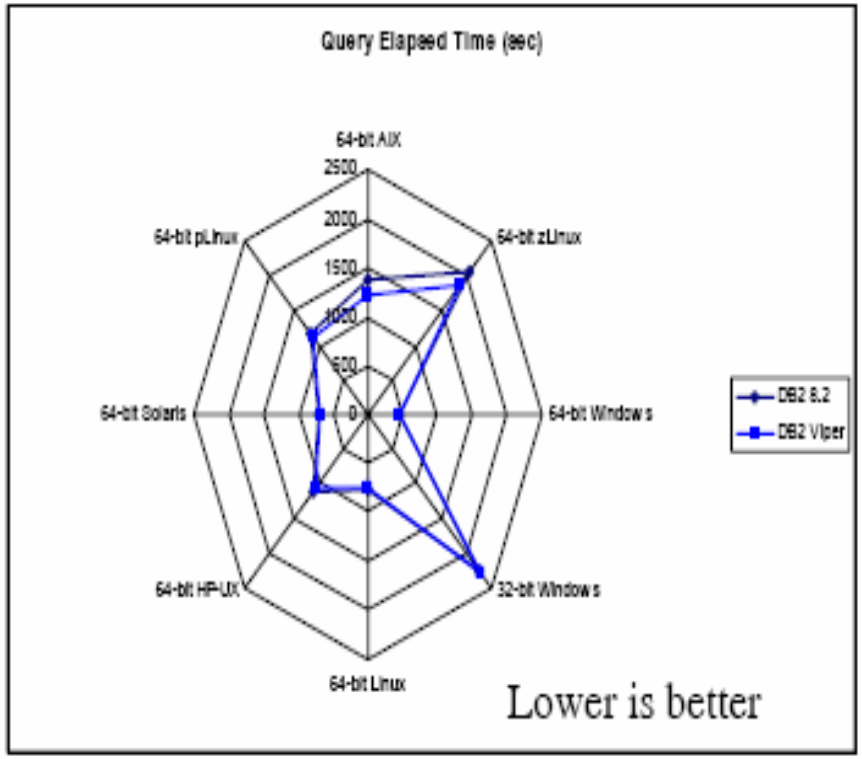

- This sample release to release comparison uses the same workload but each OS is run on different H/W i.e. Comparisons across platforms are meaningless
- No new DB2 9 capabilities are explicitly exploited

zLinux has significantly improved

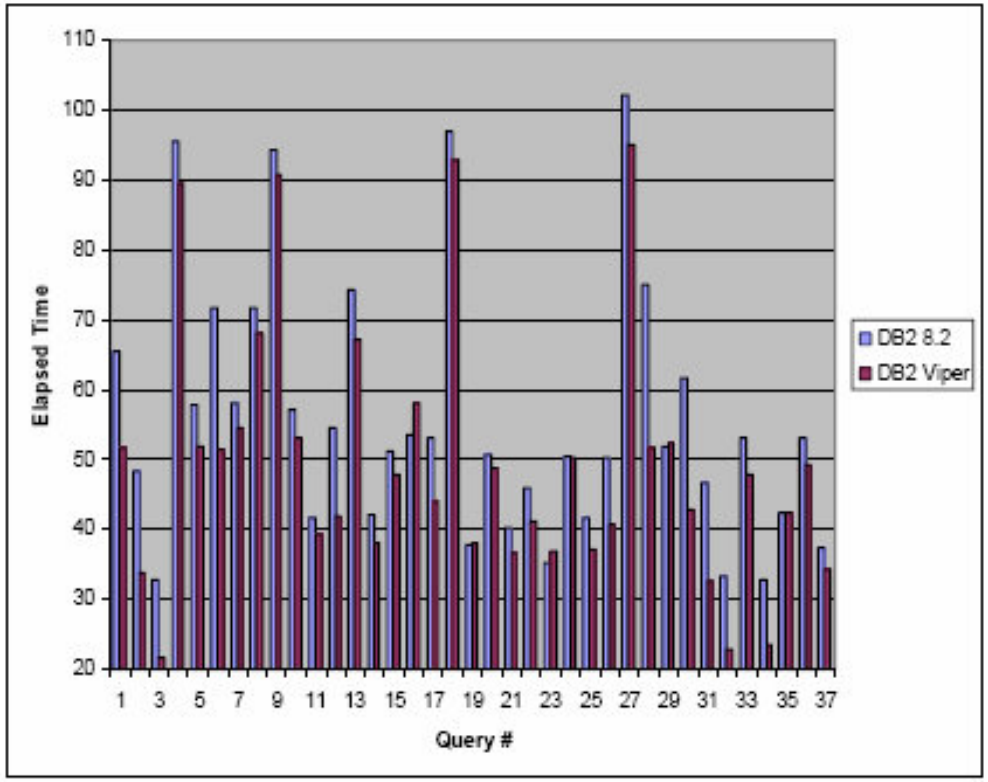

Roughly a ~13% improvement in the mean query elapsed time

IBM

Performance Characteristics of Key New DB2 9 Capability

TAKE BACK CONTROL

Compression

TAKE BACK CONTROL

Compression Results

- I/O bound system
  - DSS system, table scan
  - 43% compression
  - 43% speed-up

- **CPU bound system**
  - OLTP system, compress largest table
  - Throughput impacted by ~8%

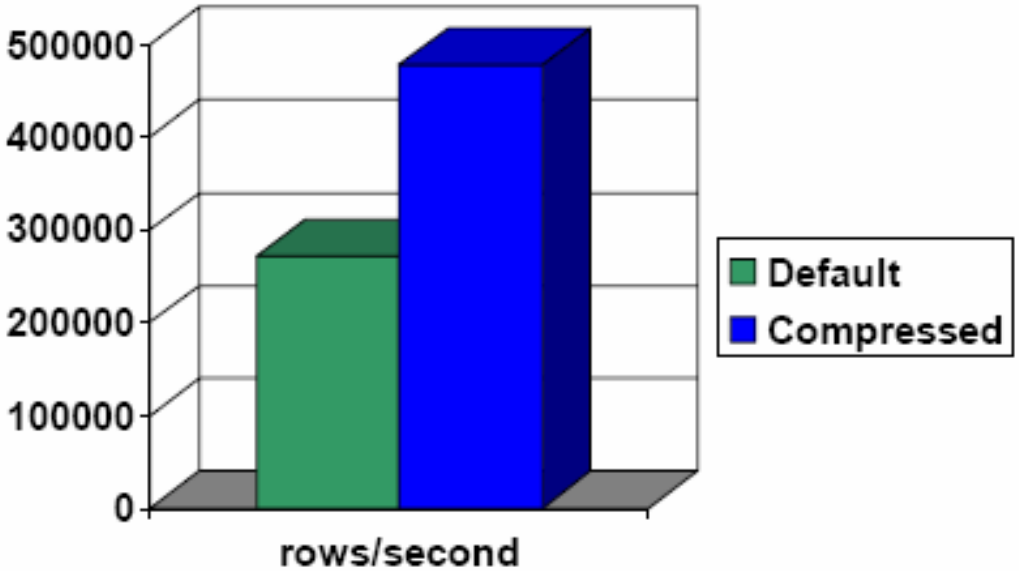

Compression Ratios at Autozone

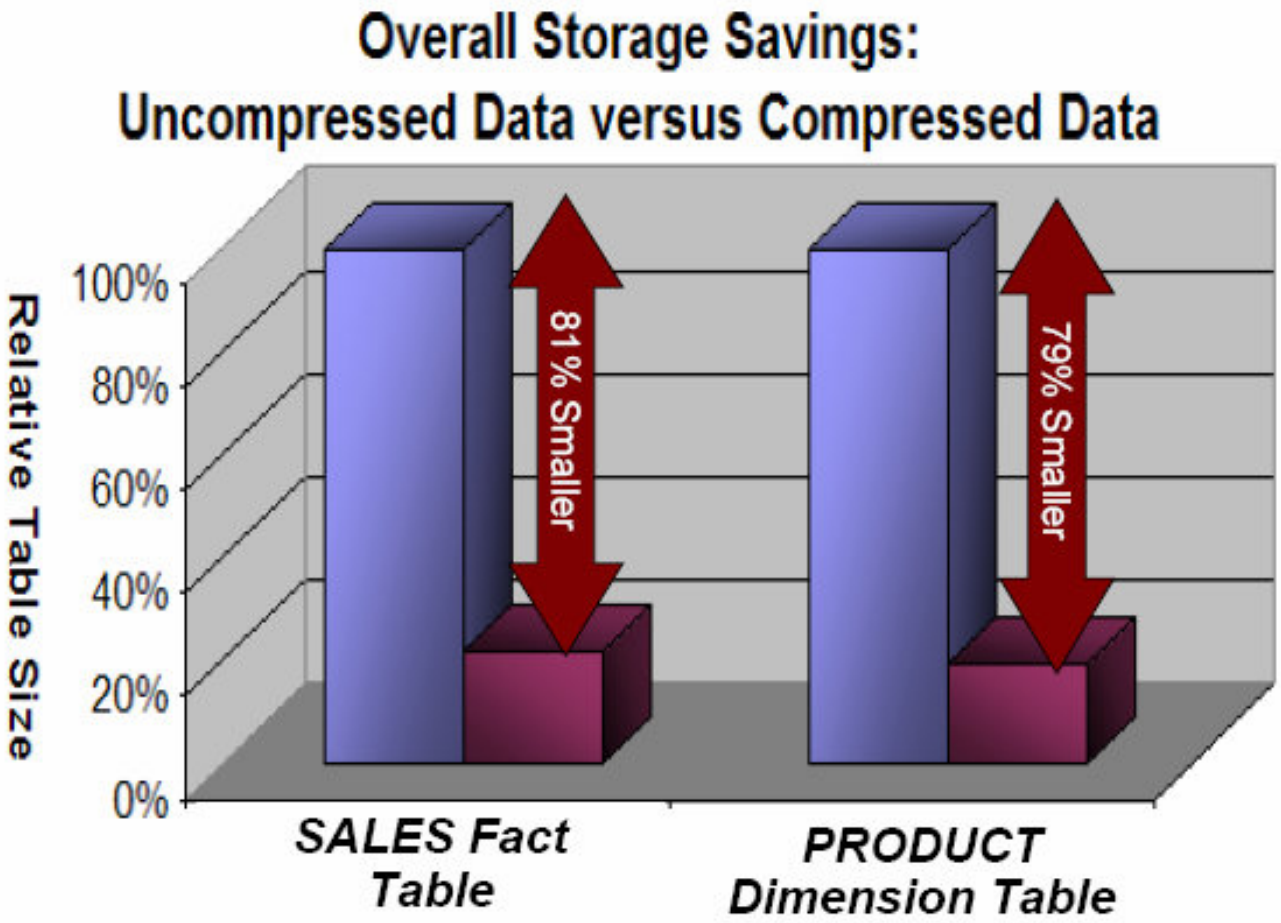

Compression Performance Gains at Autozone

Overall Workload Performance:

Uncompressed Data versus Compressed Data

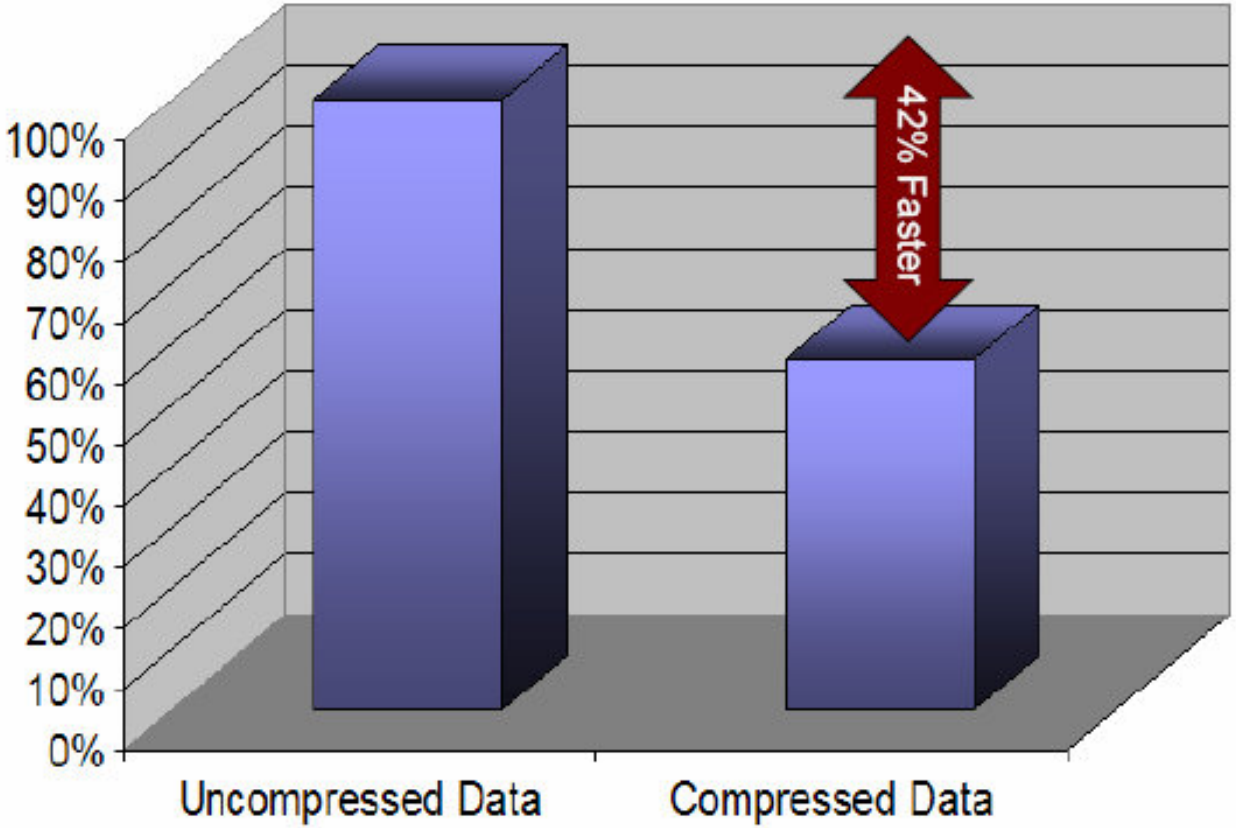

Table (Range) Partitioning

TAKE BACK CONTROL

Table Partitioning / MDC / DPF

- 3 ways to spread data (can mix and match!)
  - DISTRIBUTE BY HASH - aka hash partitioning in DPF
  - PARTITION BY RANGE – aka table partitioning
  - ORGANIZE BY DIMENSIONS – aka MDC

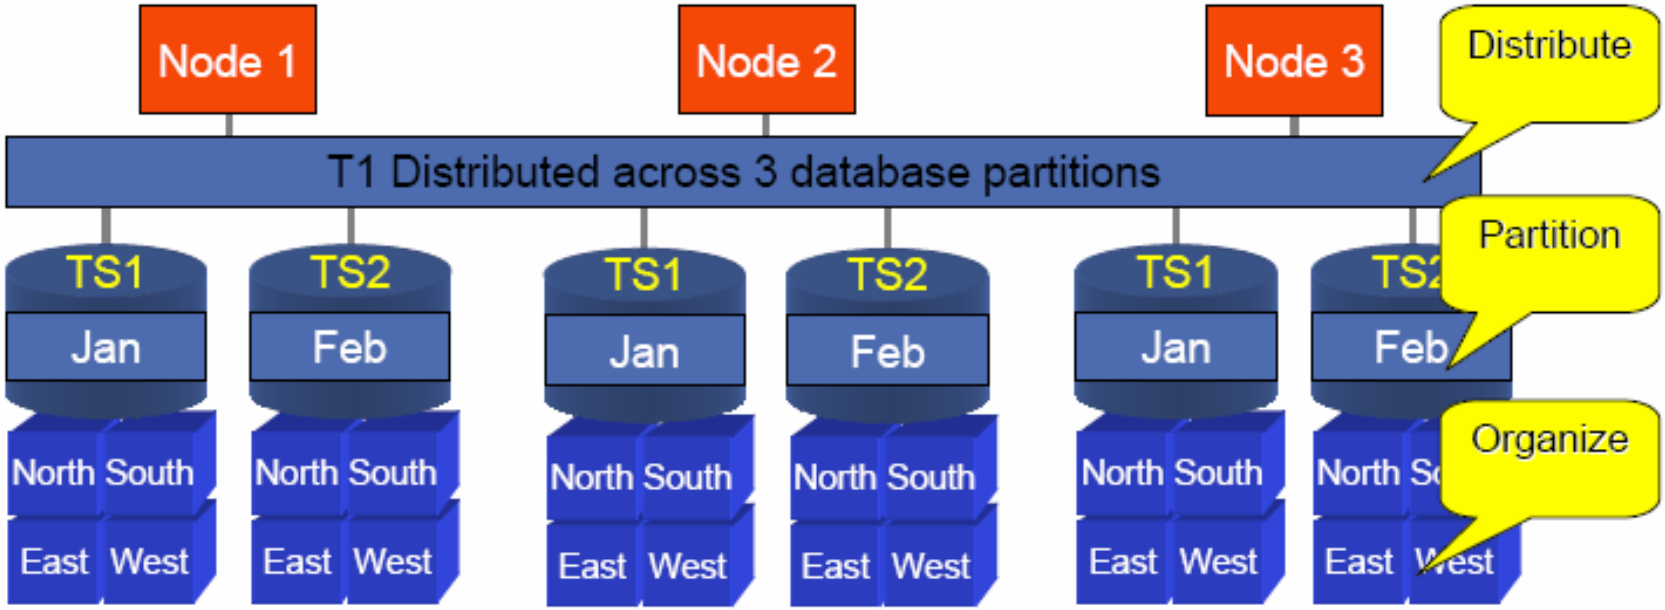

Table Partitioning and MDC

Need to Address	Best	Explanation
Data availability during roll-out	Partitioning	Use DETACH to roll-out large amounts of data with minimal disruption.
Query performance	Both	MDC best for querying multiple dimensions. Table partitioning helps via partition elimination.
Minimize need for reorg	MDC	MDC maintains clustering which reduces the need to reorg.

IBM

Autonomics

TAKE BACK CONTROL

Autonomics – what is it?

- Built-in intelligence to automate repetitive onerous DBA tasks
- DB2 9 enables significant autonomic capability by default
  - **Automatic Storage**
    - takes the fuss out of container handling
  - **Autoconfigure for db/dbm config parms**
    - performs initial basic tuning on three dozen performance parameters
  - **Auto RUNSTATs**
    - updates table and index stats every day and only if needed.
  - **Self-tuning memory manager**
    - constantly adapt memory allocation to DB2
- DB2 Design Advisor continues to offer superb recommendations

Autonomics – considerations?

- For a **fully tuned** environment
  - **STMM** overhead is ~2%
  - Aggregate **autonomic** overhead is ~3%
- For a **semi-tuned** environment
  - Benefits of STMM alone typically outweigh any overhead
- For an **untuned** environment
  - Leverage unused or under-utilized resources
  - 25% or greater improvement easily achieved

IBM

XML

TAKE BACK CONTROL

XML: A New Model

XML Developer

"I see a sophisticated XML repository that also supports SQL."

Familiar Programming Models

SQL Developer

"I see a sophisticated RDBMS that also supports XML."

Mature Services

Optimized Storage Models

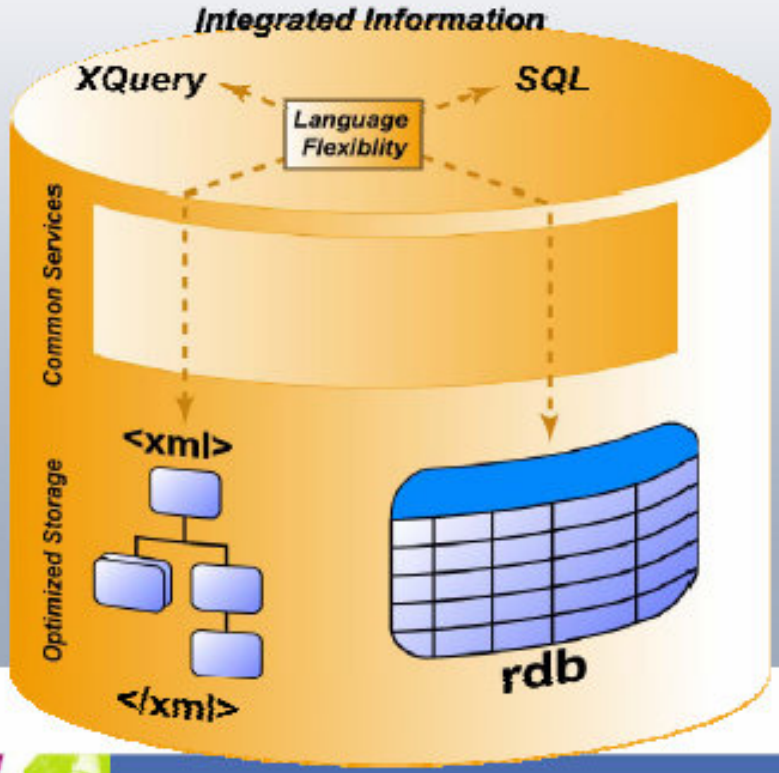

Familiar Tooling

Optimized Performance & Scale

Don’t make trade-offs – Choose DB2 Viper for XML

XML Storage Model	DB2 Multi Structured Hybrid (Relation / Hierarchical)	Oracle or Microsoft CLOB	Oracle Shredded
Schema Flexibility (0, 1, many)	**Best**	**Best**	**Bad**
Query Performance	**Best**	**Bad**	**Good**
Update Performance	**Best**	**Bad**	**Good**
Insert / Delete Performance	**Best**	**Good**	**Bad**
Full Document Retrieval	**Good**	**Best**	**Bad**
Information Fidelity	**Best**	**Best**	**Bad**
Integration w/ Relational	**Best**	**Bad**	**Good**
Subdocument level access	**Best**	**Bad**	**Good**

XML column – considerations?

- Requires a UNICODE database
- Slower to INSERT than a CLOB
  - Not surprising since it does much more
- Faster to INSERT than shredding / decomposition
- DB2 9 decomposition MUCH MUCH faster than XML extender
- Need to look at usage pattern to decide on optimal approach

Performance Proof Points
DB2/XML performance papers

- **A performance comparison of DB2 9 pureXML with CLOB and shredded XML storage**
  - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0612nicola/
- **DB2 9 pureXML Scalability on Intel® Xeon® MP Platforms Using IBM N Series Storage**
  - http://www.intel.com/cd/ids/developer/asmo-na/eng/dc/mobile/technologies/326328.htm
- **15 best practices for pureXML performance in DB2 9**
  - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0610nicola/
- **Exploit XML indexes for XML query performance in DB2**
  - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0611nicola/
- **On the Path to Efficient XML Queries**
  - http://www.vldb.org/conf/2006/p1117-balmin.pdf
- **DB2 9 XML performance characteristics**
  - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0606schiefer/

Summary

- DB2 performance leadership continues
- Core DB2 9 function is as fast or faster
- Significant new features available with DB2 9
  - Function and Performance
  - As you exploit new capabilities you will realize additional benefits
- Initial best practices are available
  - Evolution of best practices continue
- Lots of other IOD sessions from IBM and our business partners

Overall Feature Comparison

	DB2 Viper	Oracle 10gR2	Microsoft SQL Server 2005
XML	Best	Good	Good
Partitioning with MDC and DPF	Best	Good	Worst
Compression	Best	Good	Worst
Autonomics by default	Good	Worst	Best
Self Tuning Memory	Best	Good	Best
Label based access	Best	Good	Worst

DB2 for z/OS Unicode community

I am kicking off the new year by starting a DB2 for z/OS Unicode community. The purpose of the community is for customers to interact with each other on the subject of Unicode. Of course IBMers will also be involved, and I intend to moderate the community, to ensure that answers are correct and that questions get answered.

If you know any customers who would like to become involved, please let them know of the community and how they can join.

What I need to get started is an **IBM ID**. If the customers don't have one, they can get one. They are free, and you need them to look at various things on the IBM website. To get an IBM ID, you go to https://www.ibm.com/account/profile/us?page=reg and request one.

Then, the customer would send me a note with their ID, and I will auth them to the Quickplace. From there, they will get a note with a link to the quickplace, they will log on with their IBM ID, and they will be "in".

For IBMers, the processes is simpler, your intranet ID is used, so I only need a note from the IBMer, and I can do the auth.

The Quickplace is pretty simple and self explanatory.

Best Regards,

Chris Crone

Senior Technical Staff Member

DB2 UDB for z/OS Development

https://www-1.ibm.com/QuickPlace/unicode4db24zos/Main.nsf/h_Toc/ddeb70a292ccb2738525725900079cfa/?OpenDocument

Questions?

Thank-you!