



IBM Software Group

# LSU in Oslo

## DB2 for z/OS

**DB2** Information Management Software



Carsten Rasmussen  
clr@dk.ibm.com  
www.ibm.com/software/db2zos


@business on demand software

IBM Software Group | DB2 Information Management Software



## Greatest Hits: DB2 for z/OS V8

- ✓ High availability
- ✓ Scalability or very large database
- ✓ Java and the web
- ✓ Queries and data warehouses
- ✓ Migrating or porting applications
- ✓ Application packages



## Continuous Availability



- **Online Schema Evolution: database changes with ALTER instead of DROP / CREATE e.g. ADD partition**
- **System-Level Log Point Recovery**
- **Data Partitioned Secondary Indexes**
- **Improved LPL Recovery**
- **Additional online zparms**

3

## Scalability and Very Large Database

- Add partitions**
- Separate partitioning & clustering**
- Index improvements**
- 4096 Partitions**
- Rotate partitions**
- Extend columns**
- Optimization improvements**
- Memory and scale increased**



4

## Migrating and porting applications



- Multi-row INSERT, FETCH & UPDATE
- GET DIAGNOSTICS
- INSERT within SELECT
- IDENTITY Column enhancements
- SEQUENCES
- CURRENT PACKAGE PATH
- SQL Procedure Language
- Dynamic Scrollable Cursor, Common Table Expression, Recursion
- Scalar Fullselect
- Materialized Query Table
- UNICODE SQL, Multiple CCSIDs
- XML Publishing
- Long names, long statements...

## Enterprise Applications & : DB2 for z/OS



4.6 certified



8.45 certified



7.8 certified

- 64 bit virtual storage
- Unicode
- Schema evolution
- System-level backup and recovery
- Multi-row fetch & insert
- Multiple DISTINCT Clauses
- Lock Contention on volatile Tables
- Fast Retrieval of Most Recent Value
- Longer Table Names & Column Names
- Additional statistics
- Convert Column Type
- Altering CLUSTER Option
- Adding Columns to Index
- Index-only Access Path for VARCHAR
- Adding New Partitions
- Separate Clustering from Partitioning
- ...

# DB2 for z/OS Version 8 is

- ✓ **SQL OLTP Leadership:** name lengths, statements, scrolling, expressions, predicates, diagnostics, ...
- ✓ **Break through limitations:** storage, partitions, log
- ✓ **Performance enhancements:** index, materialized query tables, more efficient IO, multi-row
- ✓ **Database changes without an outage:** add partition, rotate partition, backup / restore
- ✓ **Integration**

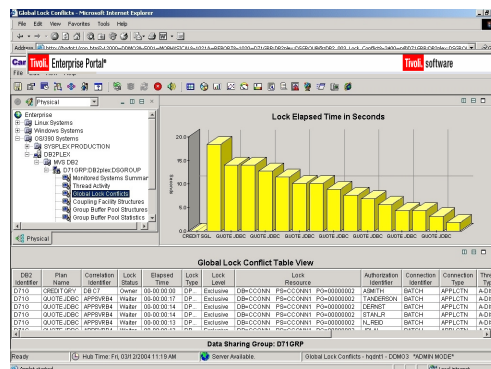
System z, z/OS & ESS platform  
 Middleware stack



*Reengineered for e-business on demand*

## Tivoli OMEGAMON XE for DB2 PE on z/OS 3.1.0

- Improved ability to monitor and manage mainframe based applications through a single integrated solution
- Familiar interfaces from DB2 PE and OMEGAMON XE products provides easy migration
  - DB2 z/OS v8 exploitation
  - DB2 Connect reporting/monitoring
  - Performance warehouse (historical data mining)
  - DB2 to CICS transaction linking
  - History monitoring
  - Event exceptions
  - Threshold checking



<http://www.ibm.com/software/tivoli/products/omegamon-xe-db2-pee-x-zos/>

## IBM System z9, z/OS & DB2 for z/OS

- ✓ **System z9 Integrated Information Processor (zIIP)**
- ✓ **Enhanced Cryptography**
- ✓ **Enhanced channels (MIDAW)**
- ✓ **Faster Processors**
- ✓ **Up to 54 Processors**
- ✓ **More memory, better value; 64 bit virtual storage**
- ✓ **z/Architecture new instructions**
- ✓ ...


Baldor Electronic



- ✓ Backup and restore
- ✓ Multilevel Security
- ✓ Unicode conversion
- ✓ Compression
- ✓ System z Application Assist Processor
- ✓ WLM enhanced
- ...



## Parallel DB2 Table Scan, EF 4K (single channel)

I/O Response Time (sec)




**Configuration:**  
 •MIDAW: z/OS 1.7  
 •Pre-MIDAW: z/OS 1.4


- DB2 for z/OS V8
- 4000 byte row size
- System z9 109
- FICON Express 2
- 2 Gbit/sec link
- DS8000 control unit




## Different Forms of Clustering



- Shared Nothing (SN) - DB2 for LUW**
- Database partitioned, no disks shared amongst nodes
  - Excellent scaling for query or partitioned OLTP
  - Distributed commit may be necessary

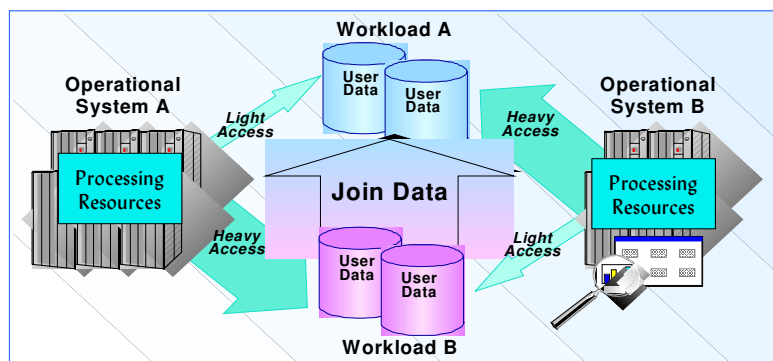


- Shared Disks (SDi) - Oracle RAC**
- No database partitioning necessary
  - But partitioning needed for good performance
  - OLTP requires inter-node concurrency and buffer coherency control mechanisms
  - Extensive messaging and disk I/O overhead
  - Dynamic load balancing not possible
  - Architecture not scalable



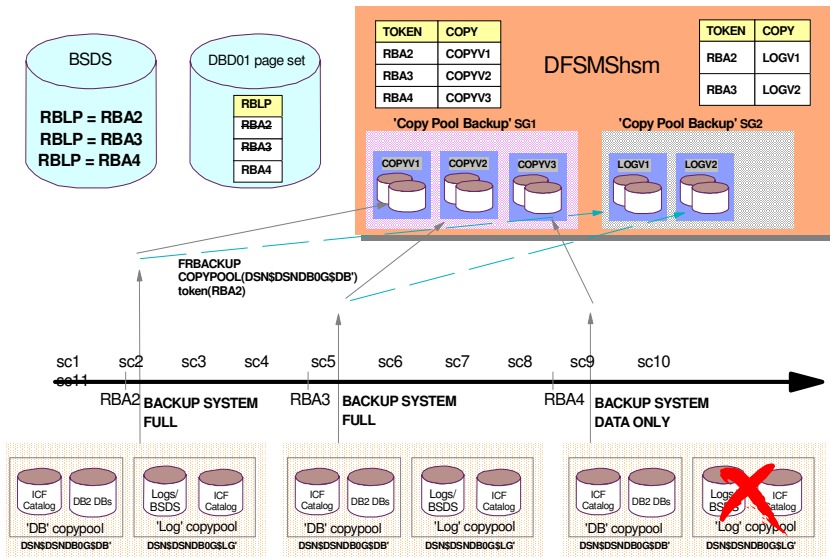
- Shared Data (SDa) - DB2 for z/OS**
- Aggressive use of z/OS Parallel SysPlex coupling facility hardware avoids heavy SDi overheads
  - Excellent scalability for OLTP, ERP and mixed workloads
  - Data partitioning not necessary for good performance
  - Strong fail-over and dynamic load balancing
  - Flexible non-disruptive growth

## Support for Mixed Workloads

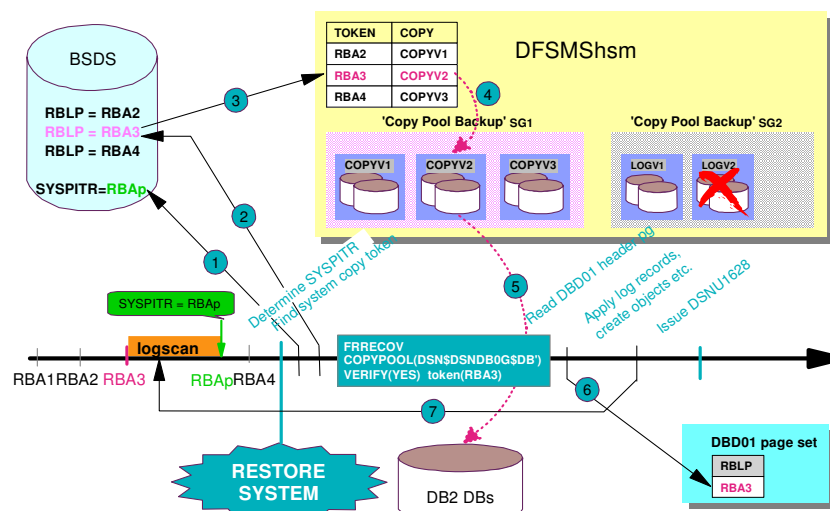


*Data Sharing*

### BACKUP SYSTEM operation



### RESTORE SYSTEM operation



Note that only the database COPYPOOL is restored


## Technology Evolution with Mainframe Specialty Engines

★ Building on a strong track record of technology innovation with specialty engines, IBM intends to introduce the System z9 Integrated Information Processor



**IBM System z9 Integrated Information Processor (IBM zIIP) planned for 2006**

Centralized data sharing across mainframes




**Internal Coupling Facility (ICF) 1997**



**Integrated Facility for Linux (IFL) 2001**

Support for new workloads and open standards




**System z9 Application Assist Processor (zAAP) 2004**

Incorporation of JAVA into existing mainframe solutions

Designed to help improve resource optimization for eligible data workloads within the enterprise

## New IBM System z9 Integrated Information Processor (IBM zIIP)

- New specialty engine for the System z9 mainframe designed to help:
  - Customers integrate data across the enterprise
  - Improve resource optimization and lower the cost of ownership for eligible data serving workloads
- z/OS manages and directs work between the general purpose processor and the zIIP
  - Number of zIIPs per z9-109 not to exceed number of standard processors. No changes to DB2 for z/OS V8 applications
- DB2 for z/OS V8 will be first user of the zIIP with
  - System z9 109
  - z/OS 1.6 or later
  - DB2 for z/OS V8
- Webcast replay [ibm.com/servers/systems/z/2006/](http://ibm.com/servers/systems/z/2006/)






## DB2 V8 and IBM zIIP can add value to database work

- Portions of the following DB2 for z/OS V8 workloads may benefit from zIIP\*:

- 1 - ERP, CRM, Business Intelligence or other enterprise applications
  - Via DRDA over a TCP/IP connection



**New Specialty Engine**

- 2 - Data warehousing applications\*
  - Requests that utilize star schema parallel queries
- 3 - DB2 for z/OS V8 utilities\*
  - Internal DB2 utility functions used to maintain index maintenance structures

\* The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP.

## Vnext DB2 Technology Themes

- Enable high-volume transaction processing for next wave of Web applications**
  - Extend the lead in transaction processing availability, scalability and performance
  - Reduce cost of ownership and zSeries-specific skill needs
  - Improve data warehousing and OLTP reporting

## DB2 SQL

z z/OS Vnext  
common

Luw Linux, Unix & Windows V8.2



- Z** { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning, TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE
- C** { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE, DELETE & MERGE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables
- L** { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE
- U**

19

## Native SQL Procedural Language

- Eliminates generated C code and compilation
- Fully integrated into the DB2 engine
- Extensive support for versioning:
  - VERSION keyword on CREATE PROCEDURE
  - CURRENT ROUTINE VERSION special register
  - ALTER ADD VERSION
  - ALTER REPLACE VERSION
  - ALTER ACTIVATE VERSION
- BIND PACKAGE with new DEPLOY keyword



20

## Decimal Floating Point

- New datatype DECFLOAT
  - Well suited to typical customer financial calculations
  - Similar to “calculator” mathematics
    - Eliminates rounding errors by using base 10 math
    - Has up to 34 digits of precision
    - Floating point convenience with fixed point precision!!!
  - Improved hardware support will be provided in the next zSeries processor generation (new IEEE standard)
    - Software emulation provided for other models




## Partition by Growth & Universal Table Space

- New partitioning scheme:
  - Single table tablespace, where each partition contains a segmented pageset (allows segmented to increase from 64GB to 16TB or 128 TB with 32K pages)
- Partition By Growth
  - Eliminates need to define partitioning key and assign key ranges
  - A new partition is created when a given partition reaches DSSIZE (defaults to 64G)
  - Retains benefits of Utilities and SQL parallelism optimizations for partitioned tables

## DB2 Vnext Themes

- ❑ Enable high-volume transaction processing for next wave of Web applications
- ❑ Extend the lead in transaction processing availability, scalability and performance
- ❑ Reduce cost of ownership and zSeries-specific skill needs
- ❑ Improve reporting

## Cost Of Ownership Trends



## Compliance/Auditing Pressure

- Regulatory compliance initiatives are impacting IT organizations in most countries/industries, and changing fast
  - Sarbanes-Oxley
  - Basel II
  - FDA: Food and Drug Administration 21 CFR Part 11
  - COPPA: Children's Online Privacy Protection Act of 2000
  - DPA: Data Protection Act (UK)
  - HIPAA: Health Insurance Portability and Accountability Act of 1996
  - PIPEDA: Personal Information Protection and Electronic Documents Act (Canada)
  - SEC Rule 17a-4: Records to be preserved by certain exchange members, brokers, dealers
  - USA Patriot Act: Uniting and Strengthening America by Providing Tools Required to Intercept and Obstruct Terrorism of 2001
- Focus is on both external threats (hackers) and internal employees

## Security in DB2 for z/OS Vnext

### Some key implementations

- Data Encryption
- Roles
- Network Trusted Contexts
- Instead of Triggers
- Improved auditing
- Secure Socket Layer



## Protecting data on disk

- We will allow encryption for the key disk resources used by DB2:
  - Tables
  - LOBs
  - Indexes
  - Image copies
  - Logs
  - Archive logs




## Auditing: DB2 Trace Filtering

- New filtering capabilities for `–START TRACE` that `INCLUDE` or `EXCLUDE` based on these keywords:
  - `USERID` -- client userid
  - `WRKSTN` -- client workstation name
  - `APPNAME` -- client application name
  - `PKGLOC` -- package `LOCATION` name
  - `PKGCOL` -- package `COLLECTION` name
  - `PKGPROG` -- `PACKAGE` name
  - `CONNID` -- connection ID
  - `CORRID` -- correlation ID
  - `ROLE` – end user's database `ROLE`

## TCO Improvements – DBA tools

- Autonomic Policy-based SQL query management/monitoring:
  - Automatic collection of performance data for long running queries
  - Automated query monitoring for the most frequent/expensive queries
  - REOPT(SMART)
  - Real time statistics exploitation by Optimizer
- Optimization Service Center (Web-based DBA admin – no 3270 screens)
  - DBA tool suite for tuning/managing SQL queries (Stats Advisor, Index Advisor, Query Rewrite Advisor, Query Workload Monitor, Resource Estimator, Query Formatter, Visual Explain, Visual Plan Hint, IBM Service Doc Generator, Partitioning/Clustering Advisor)
- Query Performance Warehouse
  - Execution history of queries
  - Identification of query patterns
  - Identification of usage patterns for tables/indexes

## DB2 for z/OS Vnext

- Integration ➤ XML, Unicode, LOBs
- Availability  
- Scalability ➤ SQL for DB2 family
- Productivity 
- ↪ Total cost of ownership ➤ Data Definition On Demand

<ftp://ftp.software.ibm.com/software/data/db2zos/VNEXT.pdf>

## Disclaimer and Trademarks

Information contained in this material has not been submitted to any formal IBM review and is distributed on "as is" basis without any warranty either expressed or implied. Measurements data have been obtained in laboratory environment. Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. The use of this information is a customer responsibility.

*IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.*

*TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, zSeries, ®*



*The FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC, WINDOWS 95*

***For additional information see [ibm.com/legal/copytrade.phtml](http://ibm.com/legal/copytrade.phtml)***