

# IBM SolutionsConnect 2013

L'IBM TechSoftware nouvelle génération  
28, 29 et 30 août - IBM Client Center Paris  
Transformer vos opportunités en succès

→ Inscrivez-vous

## DB2 11 and Beyond – Celebrating 30 Years of Superior Technology

**Maryela Weihrauch, IBM Silicon Valley Lab,  
Distinguished Engineer, DB2 for z/OS  
[weihrau@us.ibm.com](mailto:weihrau@us.ibm.com)**

**ZDB04**

# Disclaimer

*Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.*



#1




too!!



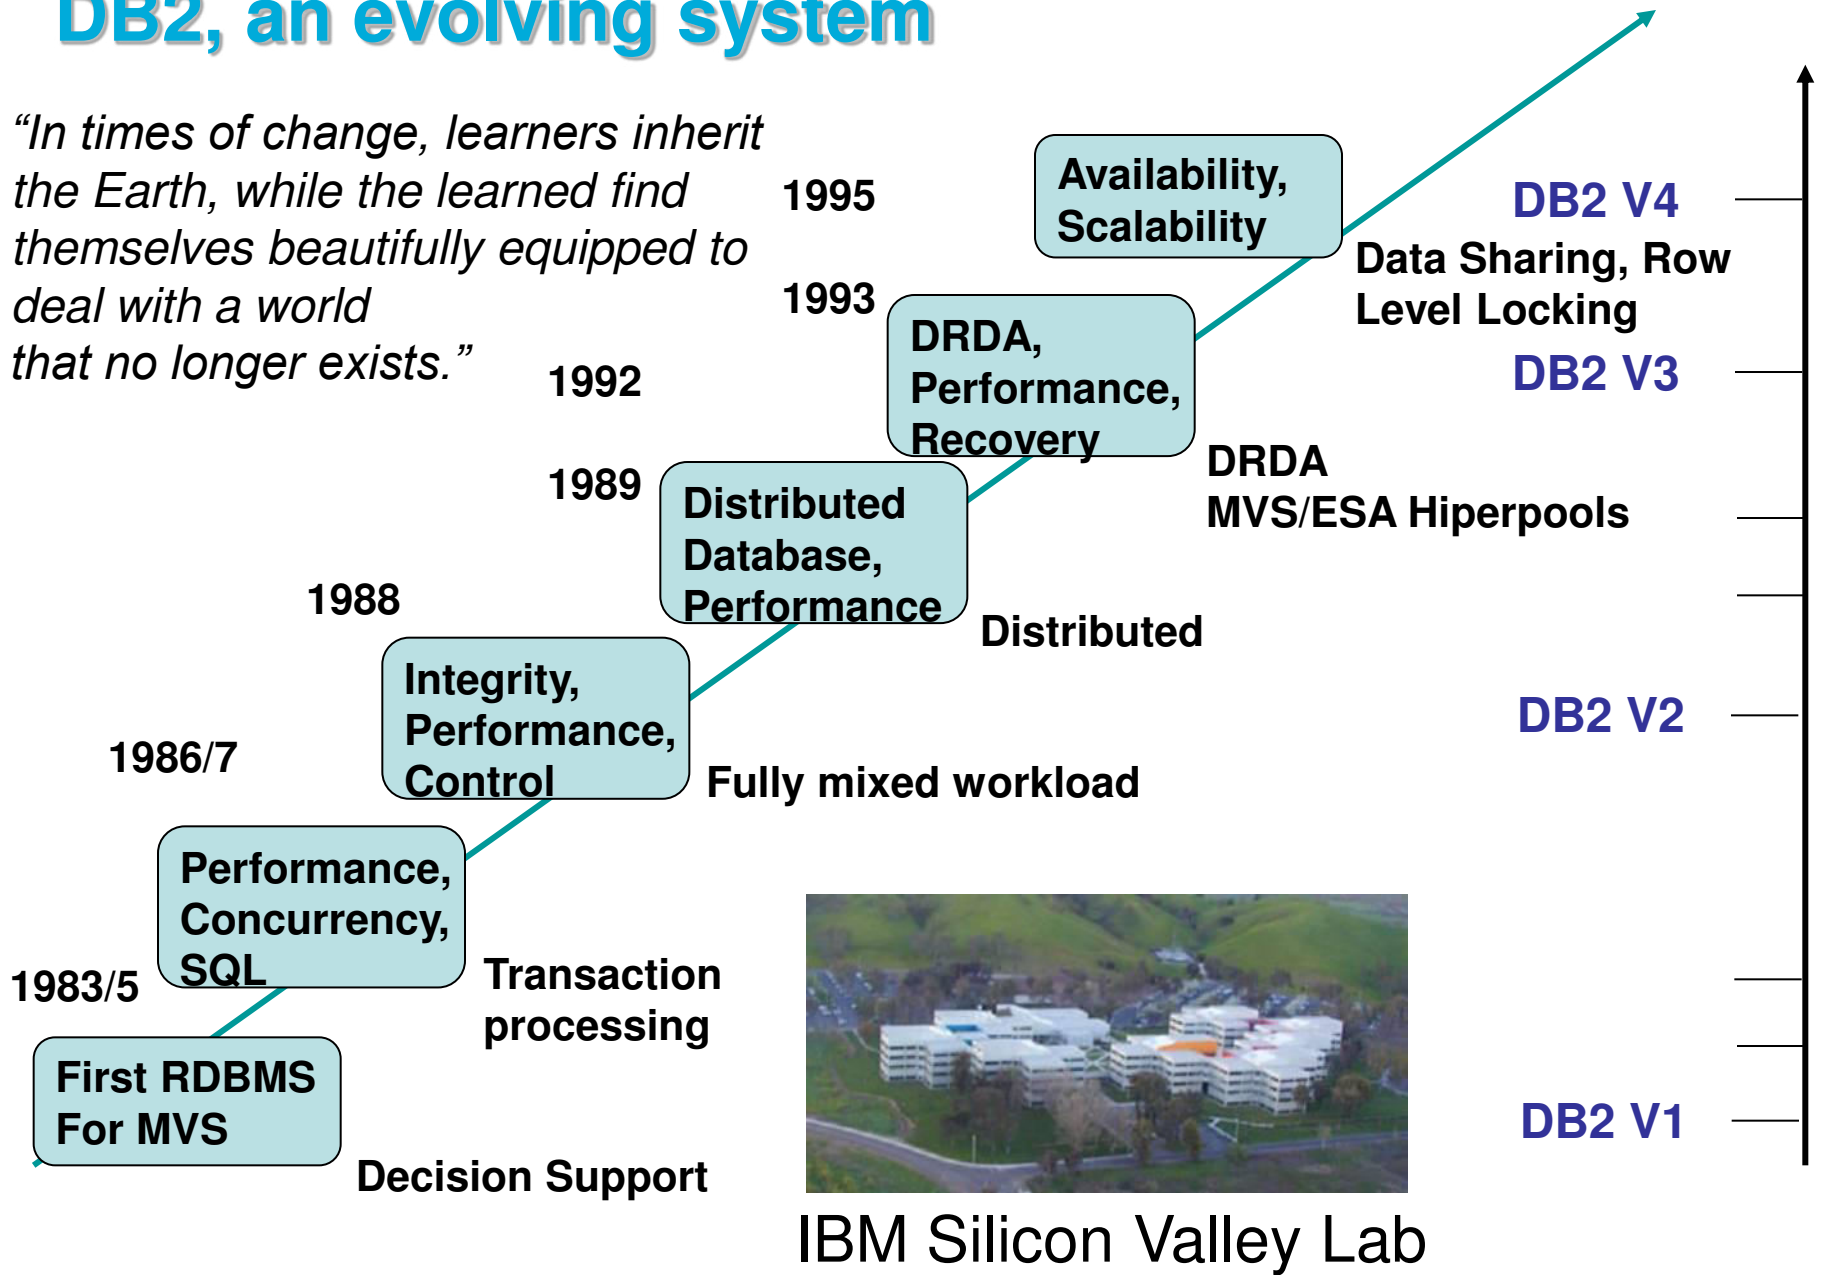
WebSphere MQ

**Dr. Edgar (Ted) Codd**  
**invented the Relational Model in 1969**  
**published in CACM 1970**

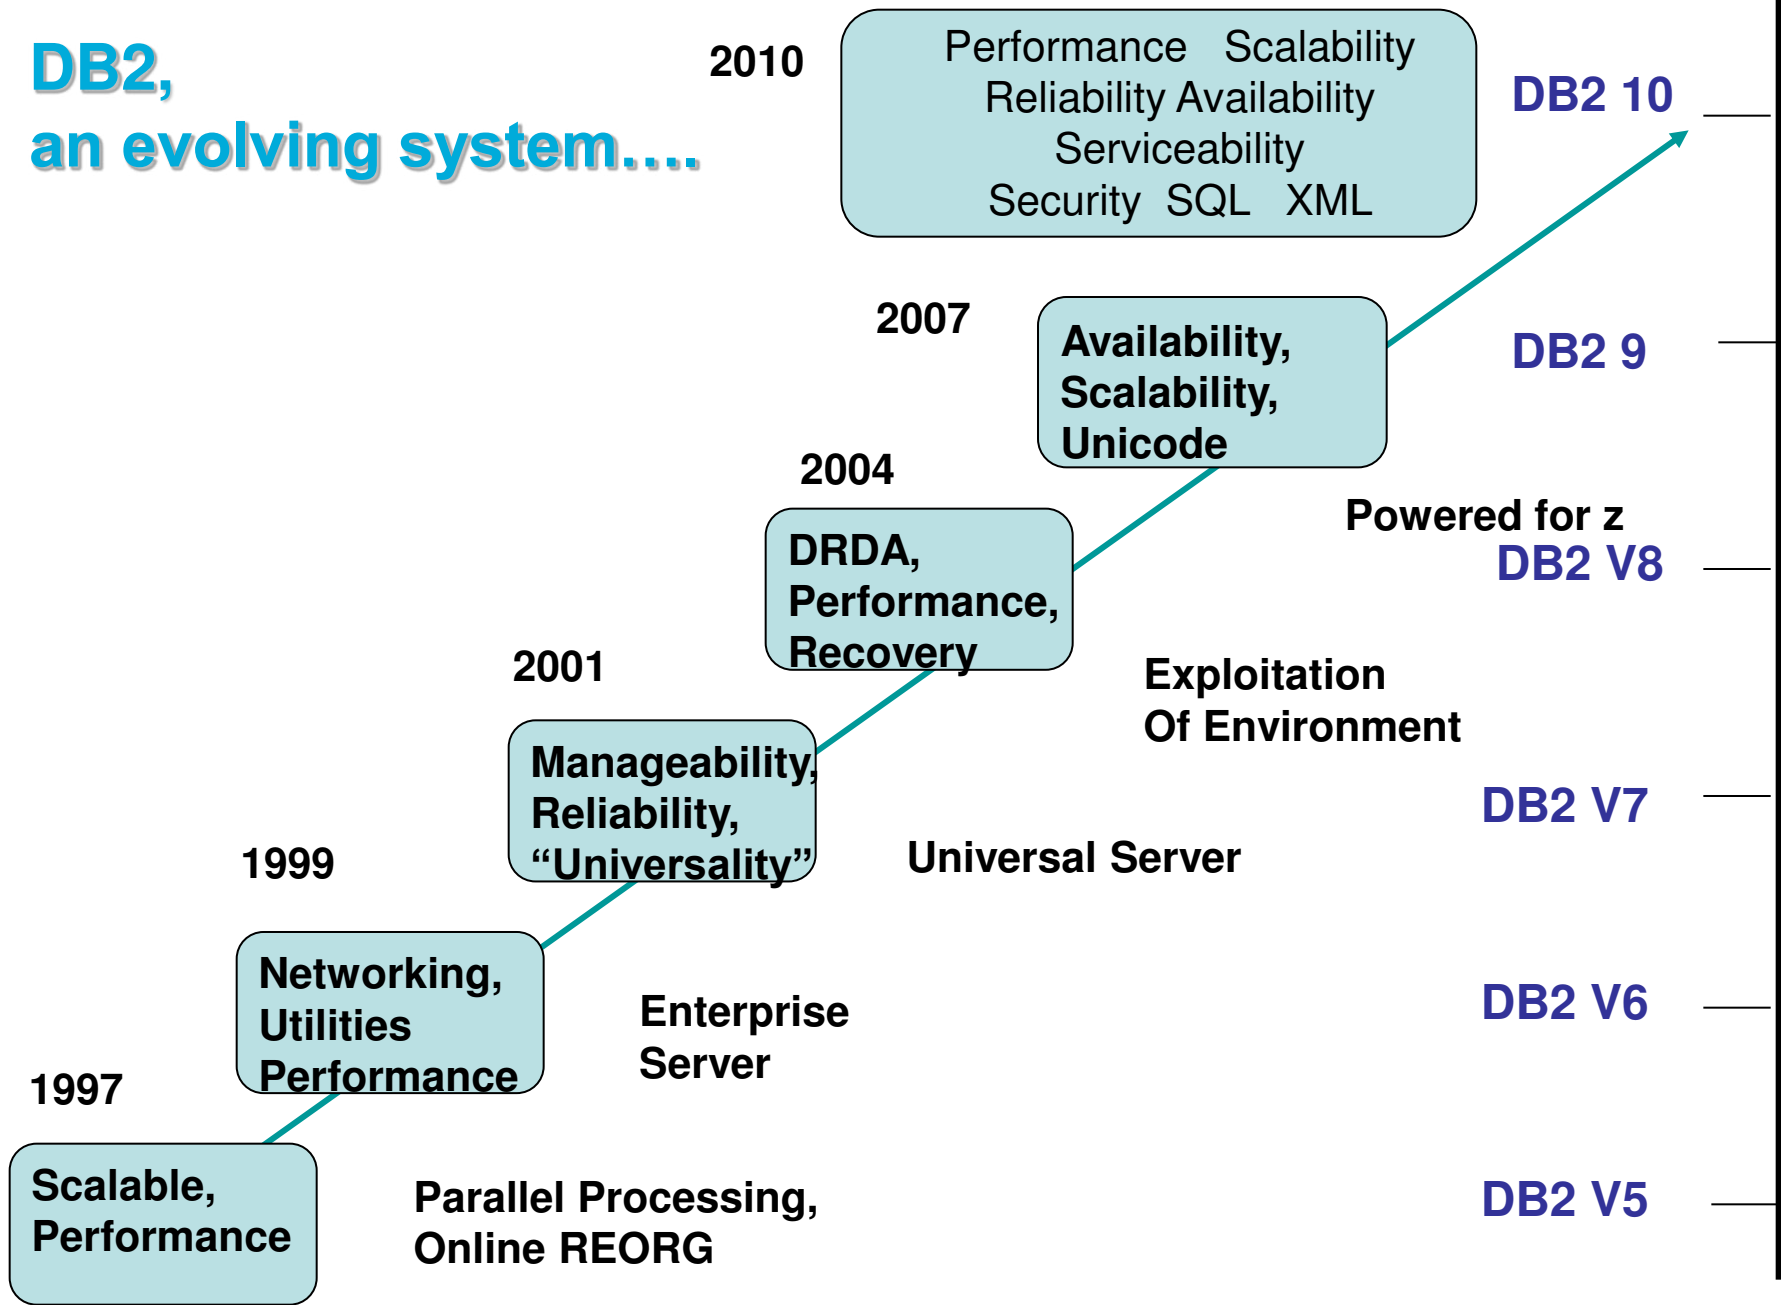


# DB2, an evolving system

*"In times of change, learners inherit the Earth, while the learned find themselves beautifully equipped to deal with a world that no longer exists."*



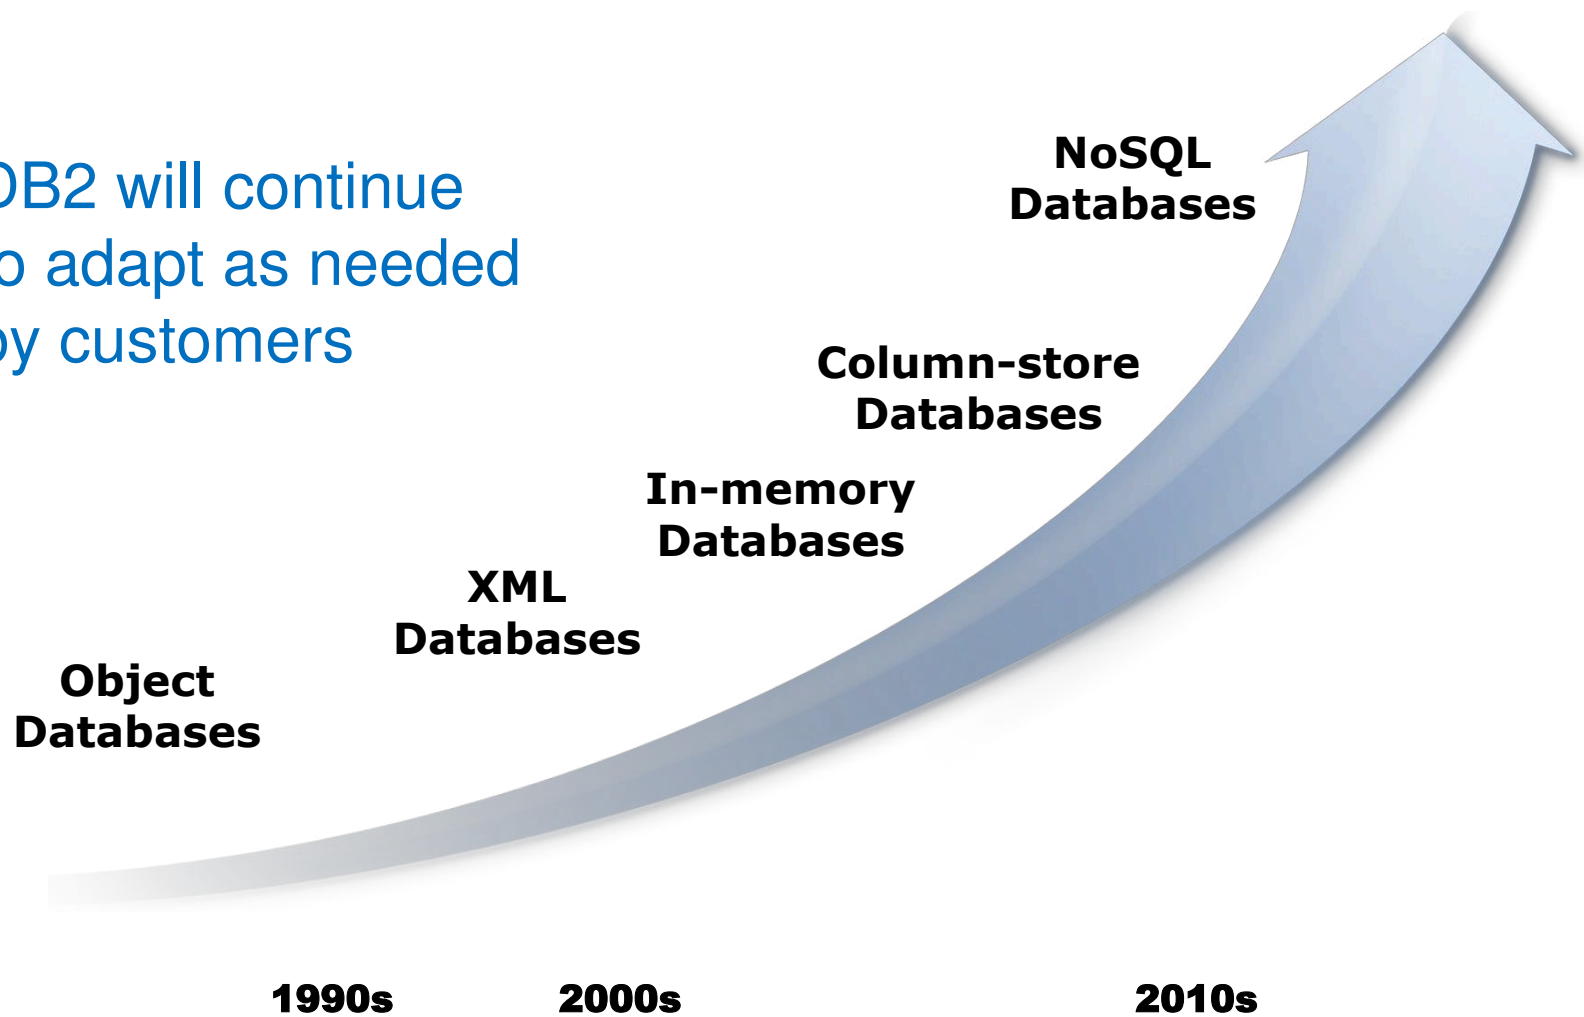
# DB2, an evolving system....





# New Technology Emerges

DB2 will continue to adapt as needed by customers



# DB2 11 Major Themes

## ■ Performance Improvements

- Improving efficiency, reducing costs, no application changes
- 0-5% for OLTP, 5-15% for update intensive batch
- 5-30% for query workloads
- Less overhead for data de-compression
- Exploitation of new zEC12 hardware features

## ■ Continuous Availability Features


- Improved autonomies which reduces costs and improves availability
- Making online changes without affecting applications
- Online REORG improvements, less disruption
- DROP COLUMN, online change of partition limit keys
- Extended log record addressing capacity (1 yottabyte)
- BIND/REBIND, DDL break into persistent threads

## ■ Enhanced business analytics

- Faster, more efficient performance for query workloads
- Temporal and SQLPL enhancements
- Transparent archiving
- SQL improvements and IDAA enhancements

## ■ Simpler, faster DB2 version upgrades

- No application changes required for DB2 upgrade
- Access path stability improvements
- Product stability: support pre GA customer production





# Performance Improvements



- **Suppress-null indexes**
  - Index entries not created when all values for indexed columns are NULL
  - Reduced index size, improved insert/update/delete performance, Oracle compatibility
  - Improved utility CREATE INDEX performance
- **INSERT performance**
  - Latch contention reduction for classes 6, 14, 19
  - CPU reduction for Insert column processing and log record creation
  - Page fix/free avoidance in GBP write
- **DDF performance improvements**
  - Reduced SRB scheduling on tcp/ip receive using new CommServer capabilities
  - Improved autocommit OLTP performance
  - DRDA package based continuous block fetch
- **zIIP enablement for all SRB-mode DB2 system agents that are not response time critical**
- **Avoid cross-memory overhead for writing log records**

# Performance Improvements...



- **ODBC/JDBC type2 performance improvements**
  - Stored procedure invocation
- **Java stored procedure multi-threading improvements**
- **RDS runtime optimizations**
  - Performance improvements for common operators
    - MOVE, CAST, output hostvar processing, CASE, SUBSTR, DATE, others
  - DECFLOAT data type performance improvements
    - Up to 23% CPU reduction for conversion to/from decfloat
    - Approx. 50% cpu reduction in INSERT, FETCH for decfloat columns
    - Helped further by zEC12 hw improvements for decimal floating point
- **Performance improvements with large number of partitions**
- **XML performance improvements**
- **Optimize RELEASE(DEALLOCATE) execution so that it is consistently better performing than RELEASE(COMMIT)**
  - Monitor # parent locks and cleanup internal structures when threshold is hit



# Performance Improvements...



- **Data compression performance improvements**
- **ACCESS DATABASE command performance**
- **DGTT performance improvements**
  - Non logged DGTTs
  - Avoid incremental binds for reduced cpu overhead
- **P-procs for LIKE predicates against Unicode tables**
- **Buffer pool improvements**
- **zEC12 exploitation:**
  - Flash memory option for buffer pools, pageable 1M frames
  - 2G page frame size
  - 1M page frames for DB2 code. Requires z/OS 2.1 or above and z/OS LFAREA
- **Latch contention reduction and other high n-way scalability improvements**

# Query Performance and Management Improvements



- **Optimizer externalization of missing statistics**
- **Query transformation improvements**
  - Convert some common stage 2 predicates to indexable (YEAR(), DATE(), SUBSTR(col,1,x), value BETWEEN COL1 AND COL2)
  - Improved indexability for OR COL IS NULL predicates
  - Enhanced pruning of "always true" and "always false" predicates
- **Index duplicate skipping**
- **Expression evaluation improvements**
- **Plan management improvements**
  - APREUSE(WARN) support
- **Selectivity overrides**
  - Improve optimizer's ability to find the cheapest access path
  - Collect filter factors for predicates in a Selectivity Profile
  - Selectivity Profile is populated via BIND QUERY



# Query Performance and Management Improvements...



- **In-memory techniques, e.g. open up sparse index, expand use of in-memory workfile**
- **RID overflow to workfile handled for Data Manager set functions**
  - DB2 10 added RID overflow to workfile
  - DB2 11 adds support for set functions (COUNT, MAX, MIN etc) which was excluded in DB2 10
- **Stage 2 predicate pushdown for RID LPF access**
  - DB2 10 added stage 2 pushdown - whereby index manager or data manager (stage 1) could call stage 2 to evaluate an expression earlier
  - DB2 11 adds this for list prefetch
- **Sort performance improvements**
- **New zparm to control max storage allocation for sort**
  - (1-128M), default=1M (same as V10)



- **Expanded RBA/LRSN. Expand to 10 bytes**
- **Increase 2G limit for a single internal DB2 storage pool**
- **BIND / DDL / Online REORG concurrency with persistent threads**
  - Use of persistent threads will increase in V10 with vstor relief
  - Examples: IMS PWFI, CICS protected entry
  - Avoid having to shut down these apps to get a REBIND through
- **DEFER DEFINE improved concurrency**
- **More online schema changes**
  - Alter partitioning limit keys
  - DROP column
  - Alter Drop Pending Changes: AREOR status is now removed
  - Point in time recovery support for deferred schema changes

- **Autonomics improvements**
  - Automatic index pseudo delete cleanup
  - Overflow row reduction
  - Optimizer externalizes missing stats to enable automated RUNSTATS
- **DDF enhanced client info fields for improved granularity**
- **New command option to externalize RTS stats (ACCESS DB)**
- **Performance monitoring improvements**
  - zIIP time added to CPU trace header
  - Package detail for rollup accounting
  - Reduction in 'not accounted for' time for query parallelism
  - Accumulated transaction summary data by connection type (new IFCID 369)
  - More granular stored procedure and UDF monitoring

- **Data Sharing Enhancements**
- **Index mgr avoid RBLDP during group restart**
- **Restart performance: fast log apply enabled**
- **SELECT from SPT01 & DBD01**
- **DESCSTAT BIND option**
- **New admin stored procedure to issue z/OS commands**
- **Compression dictionary handling for IFCID 306**

# Summary of Utilities Improvements

- **Availability**

- Online data repartitioning
  - REORG REBALANCE SHRLEVEL(CHANGE)
  - Online ALTER of limit keys
- Online REORG availability improvements
  - SWITCH phase reduction
  - Improved drain processing
- Part level inline image copies for REORG

- **Usability**

- Online REORG automated mapping tables
- Improved utility parallelism and control
- DISPLAY UTILITY enhancements

- **CPU reduction**

- More zIIP offload for LOAD and RUNSTATS

- **Performance**

- Faster LOAD processing
- Inline statistics improvements, reduced need for RUNSTATS
- Optimizer input to statistics collection
- Reduced system resources for utilities (MRU buffer management)
- DSNACCOX performance

- **Temporal data enhancements**
  - Support for views
  - Special register support
- **Transparent archive query**



# New Application Features...



- **Global variables**
  - Named memory variables that you can access and modify through SQL
  - Share relational data between SQL statements
    - Without the need for application logic to support the data transfer
- **SQLPL improvements (performance, manageability, function)**
  - Autonomous transactions
  - Array data type support
- **Alias support for Sequence objects**
  - Private alias, as currently supported for tables/views
  - Or new public alias support, enabled only for sequence objects
    - Implicit SYSPUBLIC qualifier
- **Row/Column Access Control UNION/UNION ALL support**
- **Unicode column support for an EBCDIC table (inplan 1H13)**
- **Provide REST UDFs as DB2 samples**



# New Analytics Features



- **SQL Grouping Sets, including Rollup, Cube**
  - Rollup is helpful in providing subtotalling along a hierarchical dimension such as time or geography
  - CUBE is helpful in queries that aggregate based on columns from multiple dimensions
- **IFI 306 performance improvement for CDC and IDAA v3**
- **DB2 support for IDAA V3 (rolled back to V10)**
  - Propagating DB2 changes to the accelerator as they happen
  - Detect staleness of data via RTS
  - Reducing disk storage cost by archiving data in the accelerator and maintaining the excellent performance for analytical queries: *High Performance Storage Saver*
  - Workload Manager integration and better monitoring capabilities
  - Increasing the query off-load scope via new special register CURRENT QUERY ACCELERATION
- **High performance SPSS in-database scoring via PACK/UNPACK (rolled back to v10)**
- **Hadoop access via table UDF**
  - UDFs shipped with BigInsights, uses new V11 generic table UDF capability

- **New Features**
  - Basic xQuery (retrofit to v10, PM47617, PM47618)
  - COBOL samples for XML (published on Developerworks website)
- **Features Enhancements**
  - Implicitly add doc node during insert/update
  - Crossloader support
  - Fix error reporting position predicate
  - Support xquery constructor as the source expression of insert and replace
- **Performance Enhancements**
  - Binary XML validation (*retrofit to DB2 V10*)
  - Partial validation after update
  - Date/Time Predicate Pushdown
  - XQuery(FLWOR) and XMLQUERY enhancement
  - Optimize Index Search Keys
  - XML Operator Improvements, use less storage and CPU
  - XQuery deferred construction
  - XMLTABLE pushdown cast
  - Avoid validation of validated binary XML data during LOAD

# DB2 Cypress: Early Thoughts

- **Large Table Management Improvements**
- **INSERT Performance and space management Improvements**
- **Improved query performance and management, cope better with poorly written SQL**
- **HW / SW integration and system autonomies**
- **Easier application porting, SQL improvements**
- **Expanded analytics capabilities**



Thank You