IBM Information On Demand 2007

IBM

DB2 9 for z/OS: Latest News and a Peek Into the Future

Session 1649

Jeff Josten, Distinguished Engineer, IBM Silicon Valley Lab

Act.Right.Now.

IBM INFORMATION ON DEMAND 2007
October 14 – 19, 2007
Mandalay Bay
Las Vegas, Nevada

Agenda

- V8 news
- zIIP news
- DB2 V9 highlights and news
- Futures

Version 8 News

V8 News

- SOAP UDF enhancements for authentication and security
  - APAR PK48773
- SORTNUM elimination
  - APAR PK45916
- New Zparm for relief for 254 compressed parts in LOAD/REORG
  - APAR PK51853
- V7 COBOL precompiler to allow V8 to work with old COBOL
  - APAR PK46170, PTF shipped 9/07
  - Note: V8 Co-processor does not work with old COBOL
- BatchPipes input to LOAD for fast loading
  - APAR PK34251, PTF shipped 6/07
- IDENTITYOVERRIDE keyword on LOAD utility
  - APAR PK27287, PTF available
- Allow RESTART(LIGHT) to not wait for INDOUBT URs
  - APAR PK29791, PTF available
- MQListener support for logical message grouping
  - APAR PK51290

zIIP News

DB2 & IBM zIIP add value to database work

Portions of the following DB2 for z/OS V8 and V9 workloads may benefit from zIIP (V9 in blue)*:

1 - **ERP, CRM, Business Intelligence or other enterprise applications**

- **Via DRDA over a TCP/IP connection (enclave SRBs)**
- **DB2 9 for z/OS Remote native SQL procedures**
- **DB2 9 XML parsing via DRDA to fully utilize zIIP (statement of direction)**

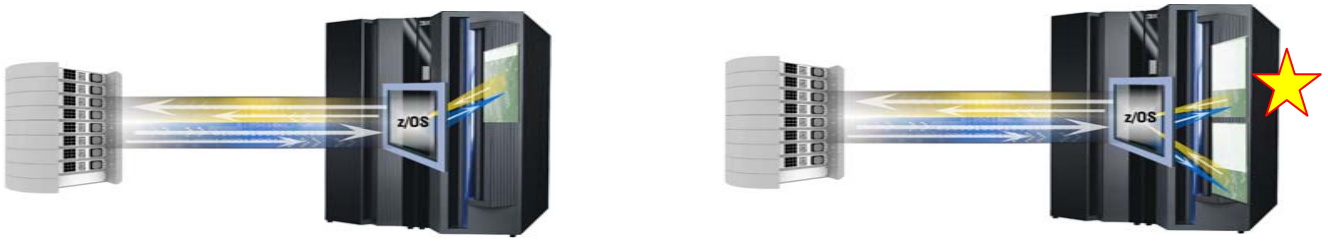

New Specialty Engine

2 - Data warehousing applications

- **Requests that use parallel queries, including star schema**
- **DB2 9 Dynamic Index ANDing, higher % of star join parallel queries eligible for zIIP**

3 - DB2 Utilities LOAD, REORG & REBUILD

- **DB2 utility functions used to maintain index structure**

*** zIIP allows a program working with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to zIIP. Above types of DB2 work are those running in enclave SRBs, of which portions can be sent to zIIP.**

V9 Highlights and News

DB2 9 for z/OS 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 9 for z/OS delivers on more than 225 requirements submitted by customers, business partners, and worldwide user group communities

64-bit Evolution and Buffer Pool Enhancements

64 bit Evolution (Virtual Storage Relief)

Virtual Storage Constraint is still an important issue for many DB2 customers. The following changes provide some relief:

- STMT CACHE BLOCKS: 100% moves above 2GB
  - Can be up to 100m or more.
- SQL STATEMENT CACHE: approx. 60% moves above 2GB
  - Statement above bar (split).
- EDM SKCT/SKPT – 100% moves above 2GB
  - This can be 100-200m for primarily static
- Other areas move above 2GB, examples:
  - Parse tree
  - Control blocks for pagesets and RTS
  - DMTR

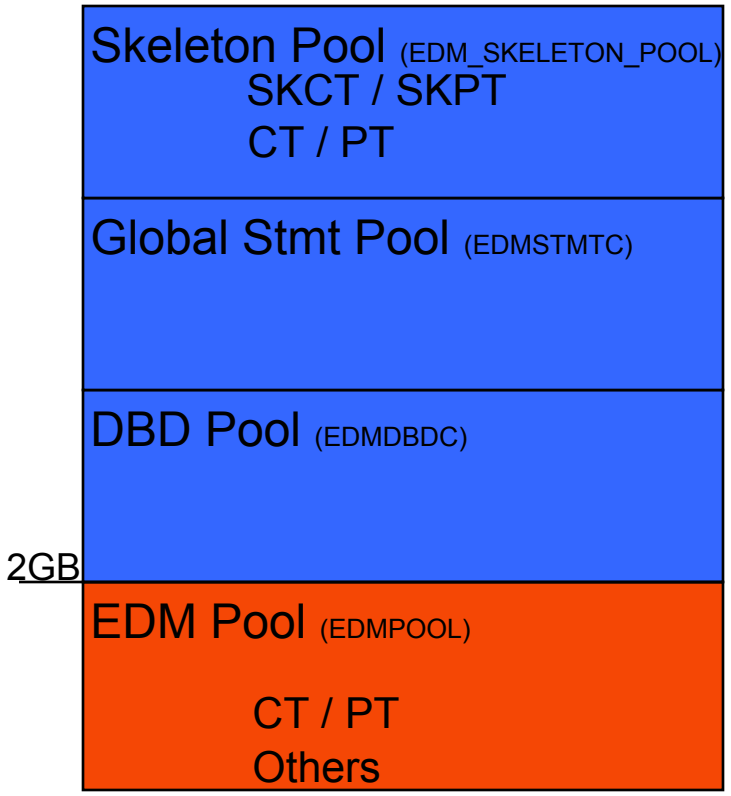

IFCID 217: detailed DBM1 virtual storage health

IFCID 225: consolidated DBM1 virtual storage health

64-bit DDF – Shared Private with DBM1

- DDF address space runs in 64-bit addressing mode
  - Shared 64-bit memory object avoids xmem moves between DBM1 and DDF and improves performance
  - VSCR

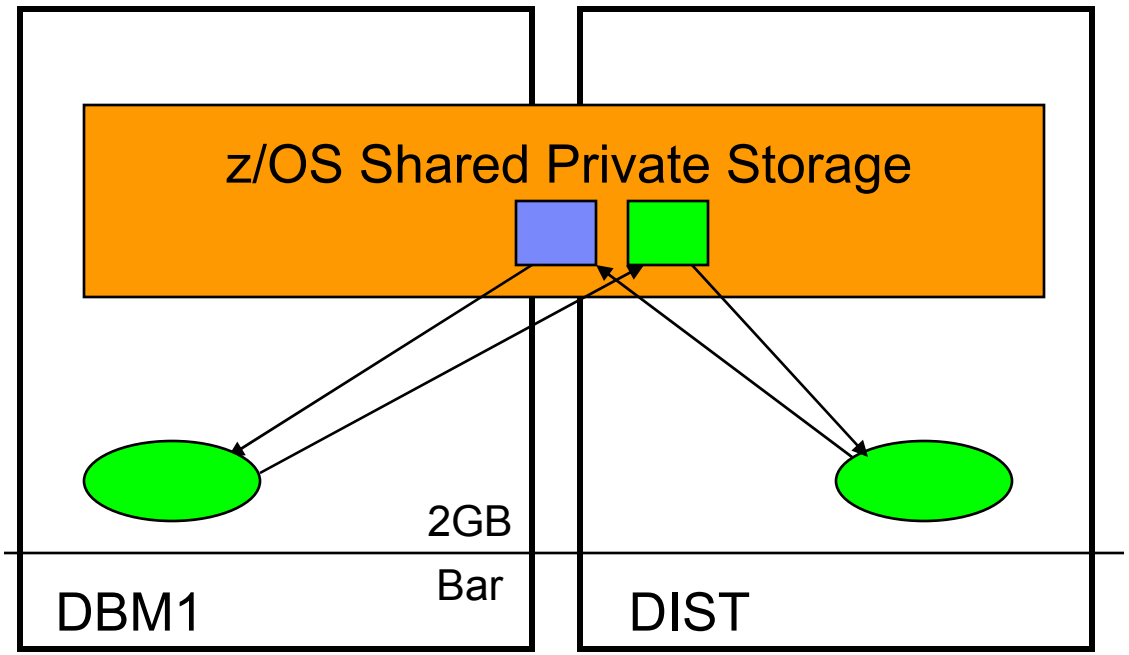

- Shared memory: new virtual storage type allowing multiple address spaces to share storage.
- Similar to ECSA – always addressable, avoids AR and XM moves.
- Different from ECSA – only available to those address spaces registering with z/OS to share this storage.
- Reduces data formatting and data movement
- Reduces virtual storage
  - It exists once, instead of in each address space

WLM Buffer Pool Management

- WLM-assisted buffer pool management
  - -ALTER BUFFERPOOL () AUTOSIZE(YES)
  - z/OS 1.8
  - DB2 registers BP to WLM and reports synch read I/O delays to WLM
  - DB2 periodically reports BP hit stats to WLM
  - WLM projects effect of adjusting BP size on workload performance goals
    - Takes into account overall system storage usage
  - WLM drives DB2 exit to adjust size if appropriate
    - V9 restricts to +/- 25%

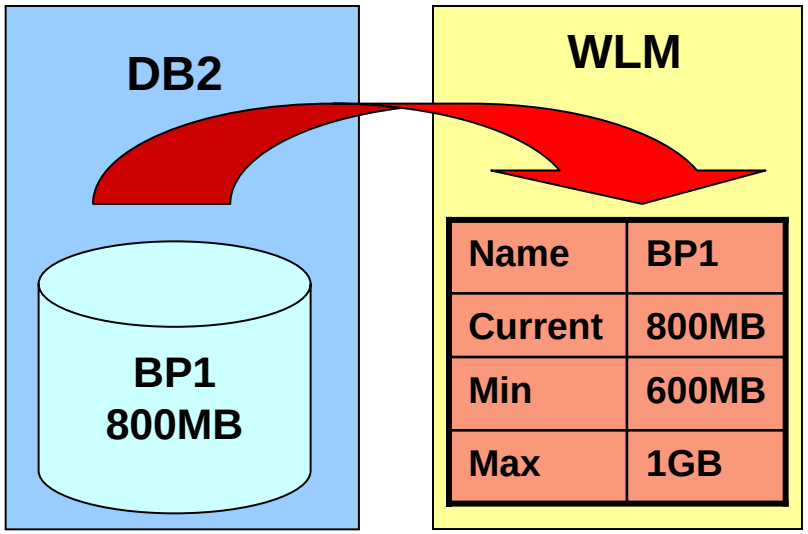

DSNB555I WLM RECOMMENDATION TO ADJUST SIZE FOR BUFFER POOL bpname HAS COMPLETED
OLD SIZE = csize BUFFERS
NEW SIZE = nsize BUFFERS

SQL Enhancements

SQL Enhancements

- INTERSECT, EXCEPT
- INSTEAD OF triggers
- MERGE
- SELECT FROM MERGE / UPDATE / DELETE
- TRUNCATE
- New data types
- New SQL functions
- Native SQL Procedures
- XML native data type

INTERSECT/EXCEPT

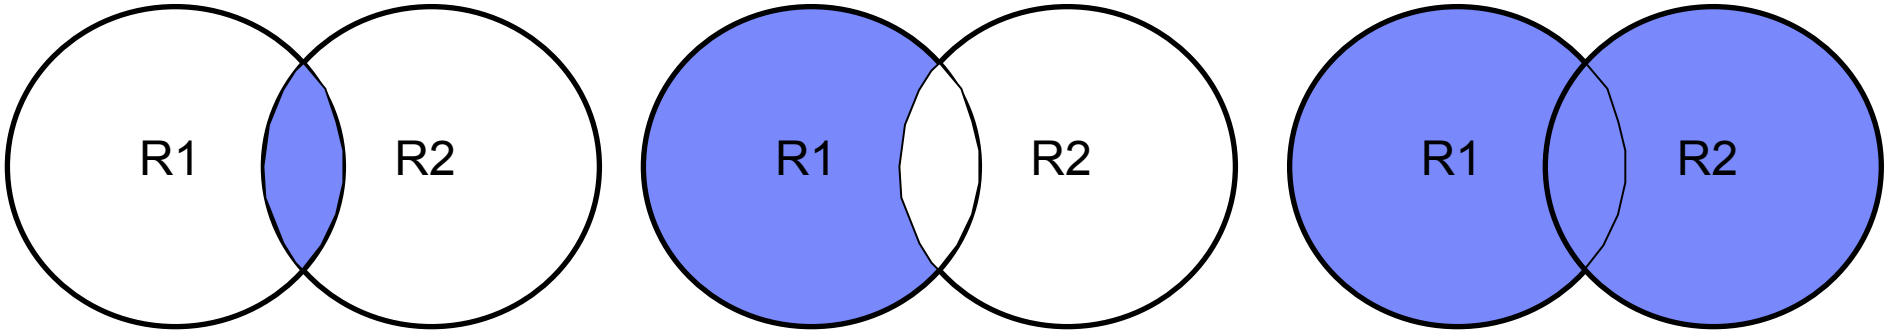

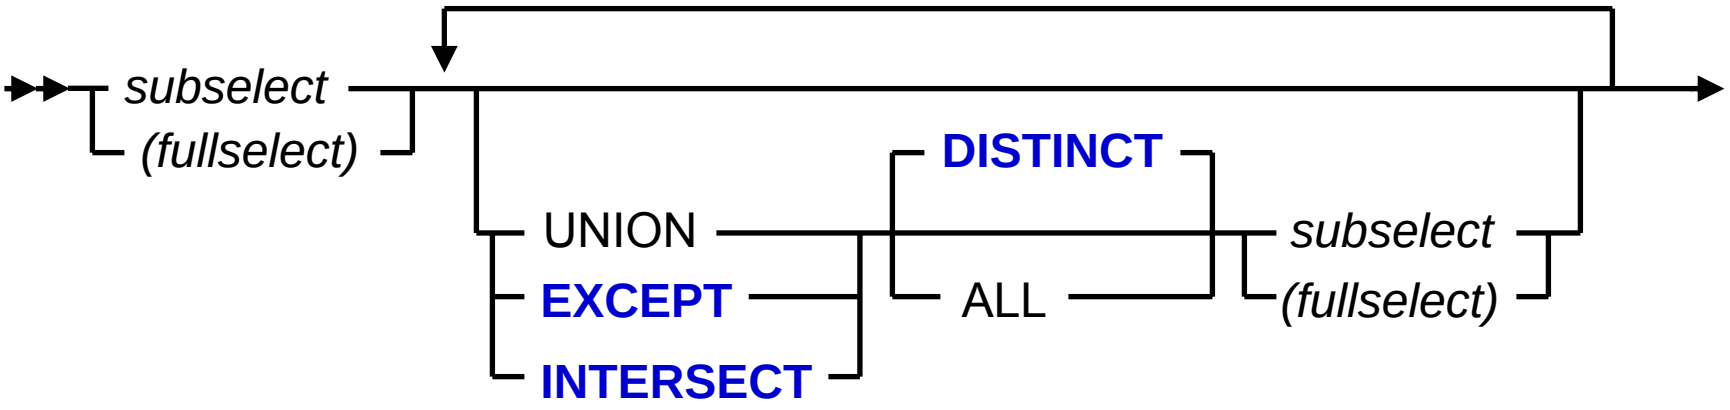

INSTEAD OF TRIGGERS

```
CREATE TABLE WEATHER (CITY VARCHAR(25), TEMPF DECIMAL(5,2))
CREATE VIEW CELCIUS_WEATHER (CITY, TEMPC) AS
      SELECT CITY, (TEMPF-32)*5.00/9.00 FROM WEATHER

CREATE TRIGGER CW_INSERT INSTEAD OF INSERT ON
        CELCIUS_WEATHER
REFERENCING NEW AS NEWCW DEFAULTS NULL
FOR EACH ROW MODE DB2SQL
    INSERT INTO WEATHER VALUES (NEWCW.CITY,
                                          9.00/5.00*NEWCW.TEMPC+32)

CREATE TRIGGER CW_UPDATE INSTEAD OF UPDATE ON
        CELCIUS_WEATHER
REFERENCING NEW AS NEWCW OLD AS OLDCW DEFAULTS NULL
FOR EACH ROW MODE DB2SQL
    UPDATE WEATHER AS W
      SET W.CITY = NEWCW.CITY,
           W.TEMPF = 9.00/5.00*NEWCW.TEMPC+32
      WHERE W.CITY = OLDCW.CITY
```


MERGE

- "Upsert"
  - A combination of insert and update, MERGE is a single SQL operation
  - Single row or multi-row
- If a row matches the ON condition it is updated,
- Otherwise it is inserted.

```
MERGE  INTO account AS T
USING VALUES (:hv_id, :hv_amt) FOR 5 ROWS AS S(id,amt)
ON T.id = S.id
WHEN MATCHED THEN
    UPDATE SET balance = T.balance + S.amt
WHEN NOT MATCHED THEN
    INSERT (id, balance) VALUES (S.id, S.amt)
NOT ATOMIC CONTINUE ON SQL EXCEPTION;
```

SELECT FROM UPDATE/DELETE/MERGE

```
SELECT athlete, score FROM FINAL TABLE
  (UPDATE events SET SCORE=SCORE+10
   WHERE EVENT='BOXING');

SELECT MEMBER_ID, UPSERT_INDICATOR FROM FINAL TABLE
  (MERGE INTO MEMBER_PROFILE AS A
     INCLUDE (UPSERT_INDICATOR CHAR(1))
     USING (VALUES (20, 'PAUL', 22) ) AS B
           (MEMBER_ID, MEMBER_NAME, MEAL_PREFERENCE)
     ON (A.MEMBER_ID = B.MEMBER_ID)
     WHEN MATCHED THEN
       UPDATE SET A.MEMBER_NAME = B.MEMBER_NAME,
                  A.MEAL_PREFERENCE = B.MEAL_PREFERENCE,
                  UPSERT_INDICATOR = 'U'
     WHEN NOT MATCHED THEN
       INSERT (MEMBER_ID,MEMBER_NAME,MEAL_PREFERENCE,UPSERT_INDICATOR)
       VALUES (B.MEMBER_ID, B.MEMBER_NAME, B.MEAL_PREFERENCE, 'I')
   NOT ATOMIC CONTINUE ON SQLEXCEPTION);
```

- First example shows SELECT returning data from the nested UPDATE
- Second example shows SELECT returning data from the nested MERGE with an INCLUDE column.
  - This include column returns the compare column from the ON clause and an indicator of whether the row was Updated or Inserted.

TRUNCATE

- Allows mass delete of all rows in base tables or declared global temporary tables
  - Simple, segmented, partitioned, universal table spaces.
  - If table contains LOB or XML columns also truncates auxiliary table spaces.
  - IMMEDIATE option – operation cannot be rolled back
    - Allows immediate reuse of allocated space for subsequent insert operations in the same UOW without committing.
- Deletes rows without firing DELETE triggers
- Option to REUSE or DROP STORAGE

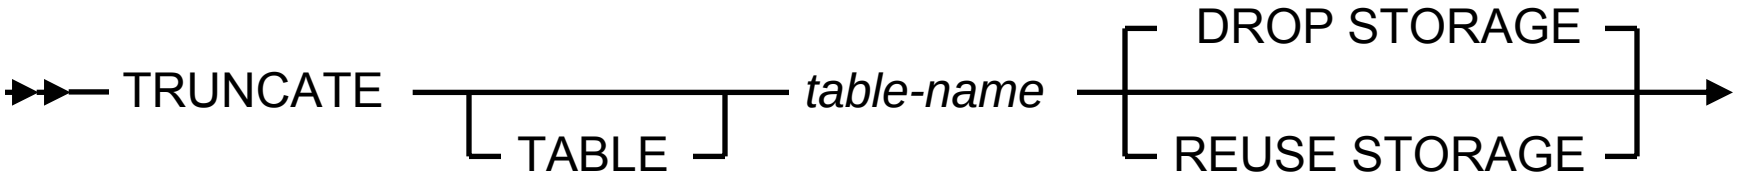

New Data Types: DECFLOAT and BIGINT

- 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
  - DECFLOAT(16)
    - 10+384 to 10-383 Positive & Negative
  - DECFLOAT(32)
    - 10+6144 to 10-6143 Positive & Negative
  - Floating point convenience with fixed point precision!
- BIGINT
  - An exact numeric capable of representing 63-bit integers
  - -9223372036854775808 to 9223372036854775807

New Data Types: BINARY & VARBINARY

- BINARY fixed-length binary string
  - 1 to 255 bytes
- VARBINARY variable-length binary string
  - 1 to 32704 bytes; maximum length determined by the maximum record size associated with the table
- Compatible with BLOBs
- Not compatible with character string data types
  - Similar to FOR BIT DATA character strings
  - Can use CAST to change FOR BIT DATA character string into binary string
- Prior to V9, FOR BIT DATA column use character based padding

Built-in Functions

- OLAP
  - RANK
  - DENSE_RANK
  - ROW_NUMBER
- Date and Time
  - TIMESTAMPADD
  - TIMESTAMP_ISO
  - TIMESTAMP_FORMAT
  - EXTRACT
  - MONTHS_BETWEEN
- Sound representation
  - SOUNDEX
  - DIFFERENCE
- String handling
  - ASCII_STR
  - EBCDIC_STR
  - UNICODE_STR
  - ASCII_CHR
  - EBCDIC_CHR
  - COLLATION_KEY
  - RPAD, LPAD
- And More …

Native SQL Procedural Language

- Eliminates generated C code and compilation
- Fully integrated into the DB2 engine
  - An SQL procedure created without FENCED or EXTERNAL is a native SQL procedure

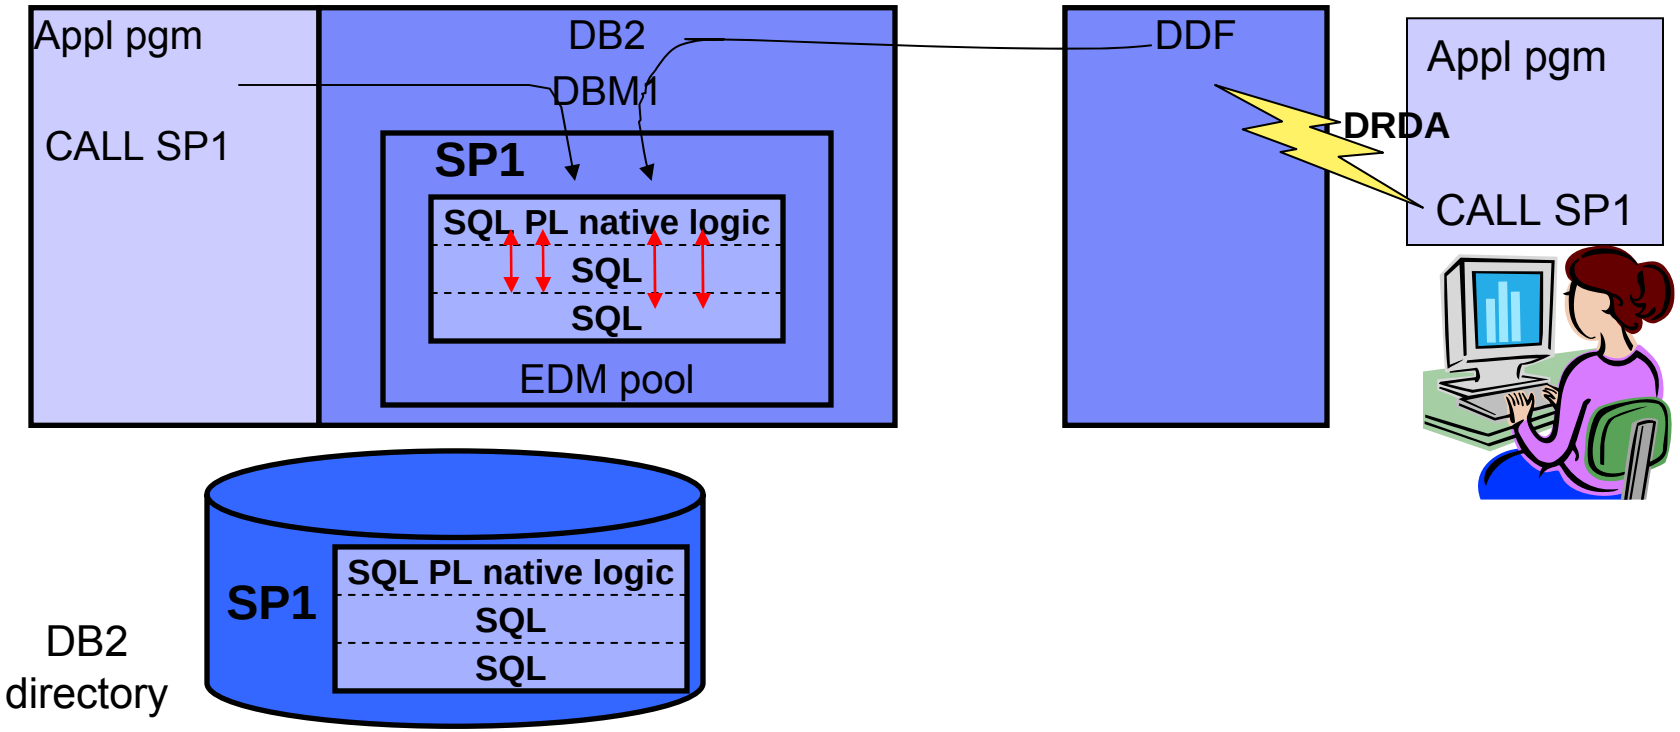

XML Capabilities Inside the Engine
pureXML™

Performance, Performance, Performance

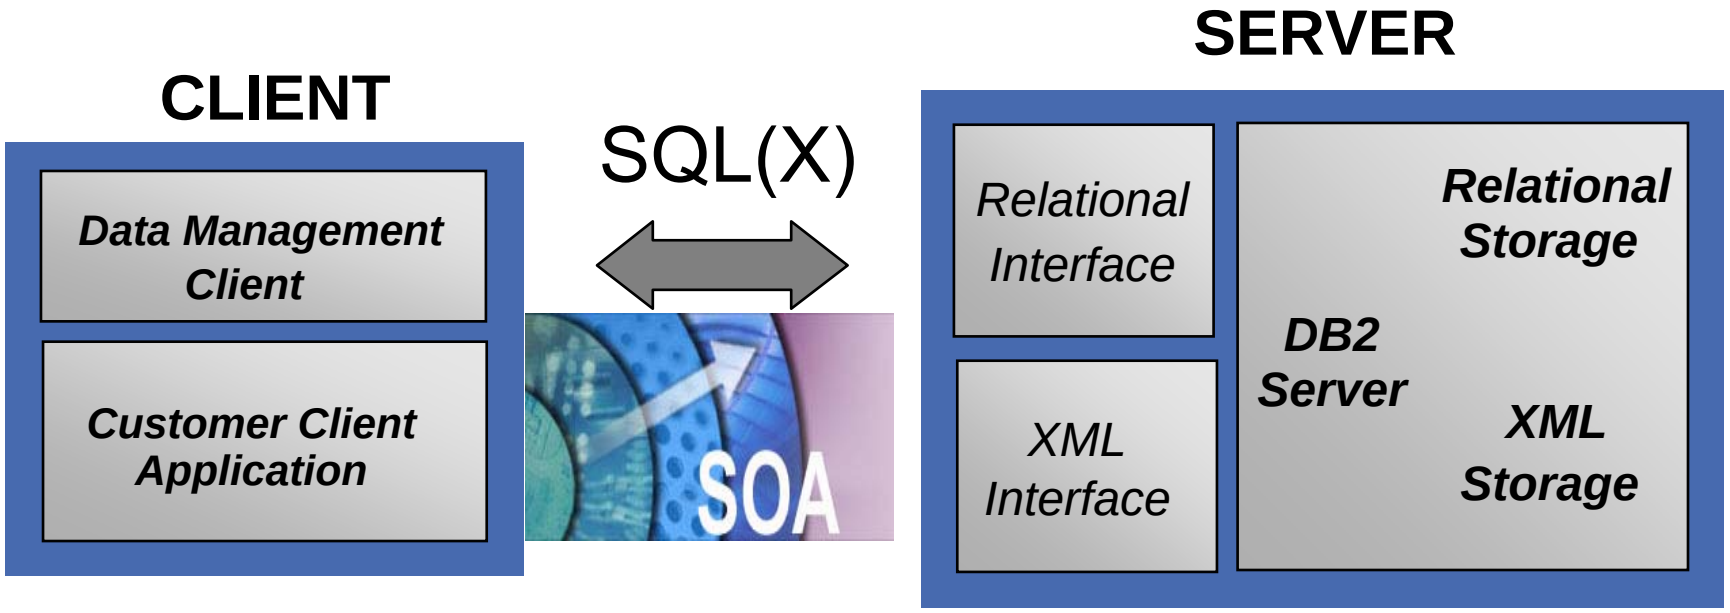

Native storage Schema Index Functions Utilities

Optimization Evolution

Dynamic Warehousing with System z
Mission-critical analysis of operational data

Rapid and secure user-access to data analysis

- *Interactive executive dashboards & information portals*

Over 50 Warehousing features in V8 and V9

- *V8 Materialized Query Tables*
- *V8 longer SQL statements, index keys, complex joins*
- *V8 up to 4096 partitions in a single table*
- *V8 and V9 Improved SQL & optimization*
- *V9 Index compression added to data compression*
- *V9 Online Rebuild Index*
- *V9 Global Query Optimization*
- *V9 Dynamic index ANDing for improved star schema query support*
- *V9 Histogram Statistics*

Cost optimization with parallel queries running on zIIP

Dynamic Index AND-ing for Star Join (Pair-Wise Join with Join Back)

- Multi-index access steps are considered independent
- Apply filtering to dimension tables before the fact table
  - Exploit single and/or multi-column fact table indexes
- Runtime assessment of filtering
  - Pre-fact dimensions with poor filter factors can be discarded at runtime and accessed post-fact
- Independent join of each dimension table to fact table via index
- The result of each pair-wise join is a set of fact table rids
- Perform Rid Sort and Rid Merge (ANDing) to generate final fact table rid list
- Final Rid list then used to retrieve data from Fact table
- Join back to dimension table(s) as necessary.

Histogram Statistics - RUNSTATS

- V8 – DB2 has data skew awareness for single values
- Histogram statistics addresses skews across ranges of data values
- Summarizes data distribution on an interval scale
- DB2 uses equal-depth histograms
  - Each quantile has about the same number of rows
  - Example - 1, 3, 3, 4, 4, 6, 7, 8, 9, 10, 12, 15 (sequenced), cut into 3 quantiles

Seq No	Low Value	High Value	Cardinality	Frequency
1	1	4	3	5/12
2	6	9	4	4/12
3	10	15	3	3/12

Security Enhancements

Trusted Contexts – An Introduction

- A ***TRUSTED CONTEXT*** establishes a trusted relationship between DB2 and an external entity such as a middleware server. For example:
  - WebSphere Application Server
  - Lotus Domino
  - SAP NetWeaver
  - PeopleSoft V7

- A set of *trust attributes* is evaluated to determine if a specific context is to be trusted.
- A trusted context allows the external entity to use a database connection under a different user ID without the database server authenticating that ID.
- It also allows an AUTHID to acquire database privileges associated with that trusted context, and not available outside it, via a ***ROLE***.

Roles and Context-specific Privileges

- **Roles** provide the flexibility to grant privileges to an AUTHID only when the user is connected via a trusted context.
- They greatly simplify management of authorization.
- An individual **role** can be defined for any AUTHID using the trusted connection, in which case the user inherits the privileges granted to the individual **role**.
- Where there is no individual **role**, any AUTHID using a trusted context inherits the privileges of the trusted context's default **role**, if defined.

```
CREATE TRUSTED CONTEXT CTX1
BASED UPON CONNECTION USING SYSTEM AUTHID WASADM1
DEFAULT ROLE CTXROLE
ATTRIBUTES (ADDRESS '9.67.40.219')
ENABLE
WITH USE FOR JOE ROLE JROLE;
```

Application Design

Universal Table Spaces

- **Universal Table Space**
  - Combination of segmented with partitioning options
    - Better space management
    - Support of mass deletes / TRUNCATE
  - If partitioned
    - Still must be one table per table space
    - Can choose Range Based partitioning (as before: PBR)
    - Can choose Partitioned By Growth (PBG)
  - DROP / CREATE to migrate existing page sets
  - Simple table spaces can not be created
    - Default table space is now Segmented

Other Miscellaneous Performance Improvements

- Reordered Row Format

Prefix	Fixed Length Cols	Varchar Indicators	Varying Length Cols

- Table APPEND option for fast INSERT
- Archive logs enabled for I/O striping

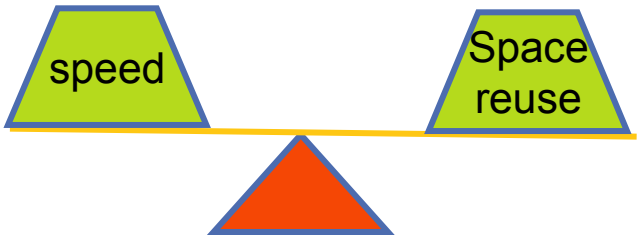

CLONE Tables

- Allows fast replacing production data without renames and rebinds
  - An alternative to online load replace
- ALTER TABLE to create a Clone Table
  - All indexes, LOB and XML objects are also cloned
  - Structures ONLY – not data
  - Base and Clone tables share the same table space and index names
  - Underlying data sets are differentiated by a data set instance number
- On single-table table spaces (partitioned or non-partitioned)
- Use insert or load to populate clone tables
- Utilities (except RUNSTATS) can operate on clone tables with new CLONE keyword
- EXCHANGE DATA switches logical names with underlying data

NOT LOGGED Tables

- Is actually NOT LOGGED tables, indexes, LOB, XML
- ALTER / CREATE a TABLESPACE as NOT LOGGED
  - ALTER not allowed if in same UOW with an update to the table space
- Indexes, LOB, and XML inherent the logging attribute of the base
  - This are considered "*Linked*" objects
- Effects the UNDO / REDO records
  - Control information is still logged
- LOB continue to log system pages & auxiliary indexes
- A Unit of Recovery (UR) is still created

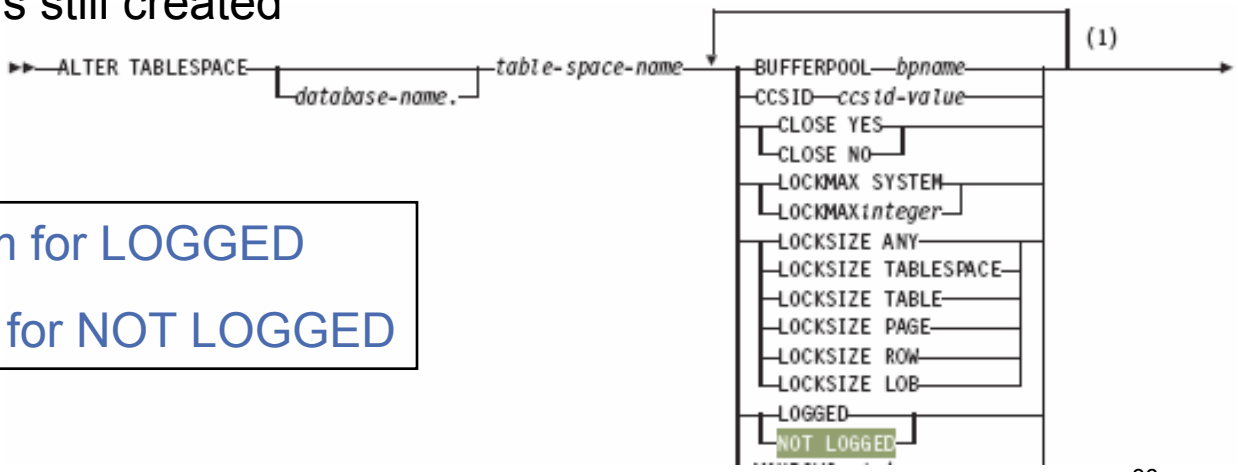

LOG YES is a synonym for LOGGED

LOG NO is a synonym for NOT LOGGED

Other Design Changes

- Optimistic Locking Support
  - Built-in timestamp for each row or page
    - Automatically updated by DB2 as a GENERATED ALWAYS column
    - Allows timestamp predicate to validate that row has not changed since last access
- Implicit Database Support
  - Automatic DB & TS
  - MAXPARTITIONS defaults to 256
  - Support for Primary Key and Unique Keys, LOB and XML objects
- STOGROUP support for SMS classes
- ALTER TABLE RENAME Column
- RENAME INDEX
- New CATMAINT options
  - Switch schema name
  - Change from owner to role (NFM)
  - Change VCAT
- RTS enhancement to identify unused indexes

LOB Improvements

- Progressive Streaming for LOB Locator Values
  - DB2 uses LOB size to determine whether to send LOB data to Java or DB2 CLI clients in one go (<32KB), in chunks (<1MB) or as LOB locator (>=1MB)
    - Transparent to application using LOB locators
- FETCH CONTINUE
  - Allows applications to retrieve LOB/XML data in pieces without the use of locators
- File reference variables
  - A file reference variable allows direct transfer of LOB data between DB2 and the file named in the variable
- Utility Changes
  - LOAD / Cross load LOB column lengths > 32KB supported
  - Logging for > 1GB LOBs
  - REORG LOB reclaim space
  - Online CHECK LOB and DATA
- Elimination of LOB locks for improved availability and performance

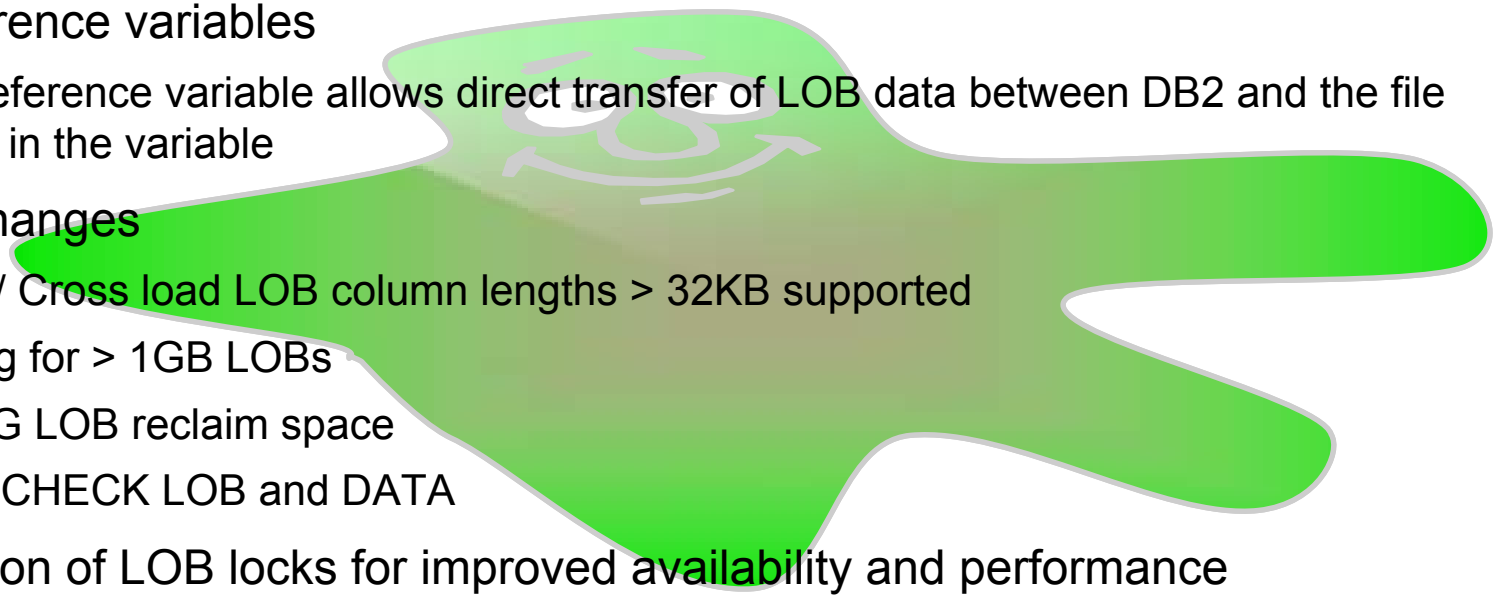

Text Search Server – Planned for V9 via upcoming Accessories Suite v1.2

- Text search for CHAR, VARCHAR, CLOB & XML columns
- Provide a text index server
- Efficient communication interaction with DB2 for z/OS
- Text indexes are persisted into DB2 tables for backup & recovery purposes

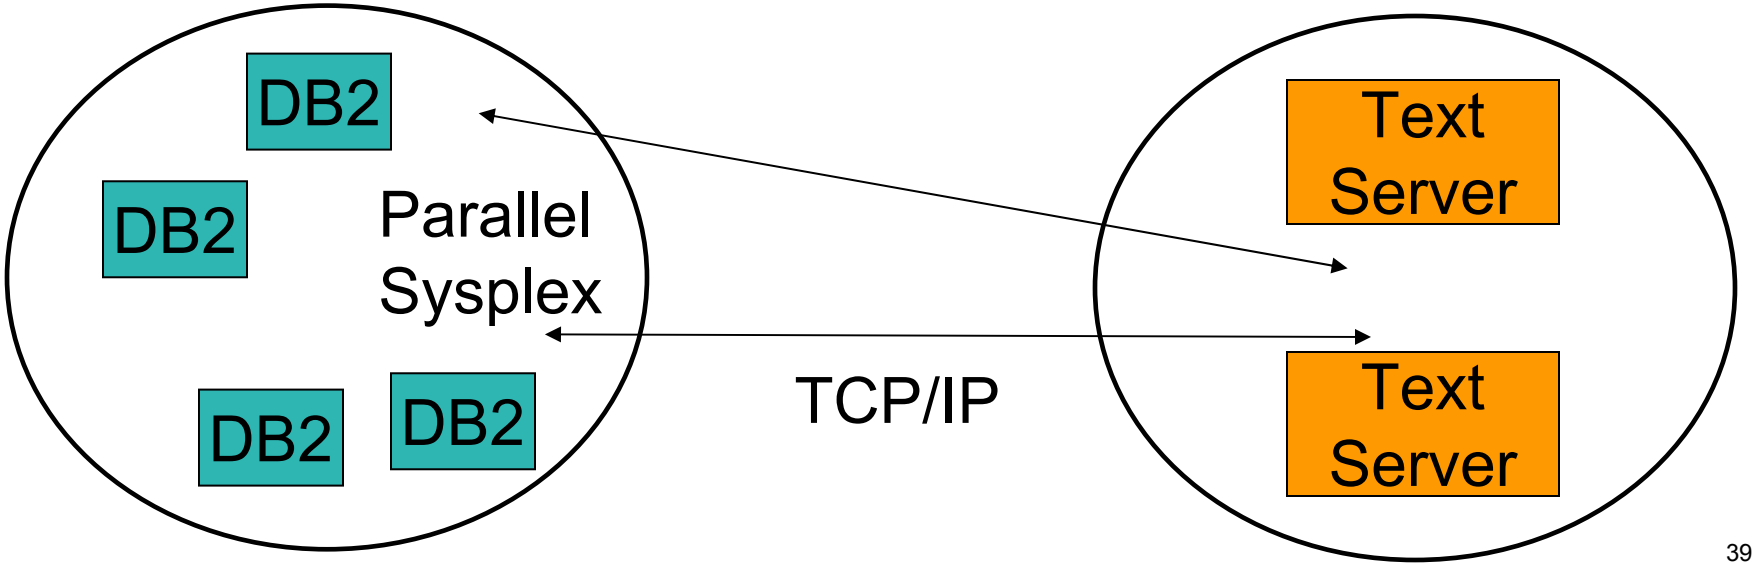

DB2 9 Spatial Support

Enabling Open Geospatial Consortium (OGC) compliant geospatial applications

- **Spatial data types**
- **Spatial functions and predicates**
- **Spatial indexes**
- **Spatial search**
- **OGC-compliant spatial catalog**

Imminent Post GA delivery for Phase II

Map

GPS

Address

Street

State

…

Indexing Enhancements

Indexing Enhancements

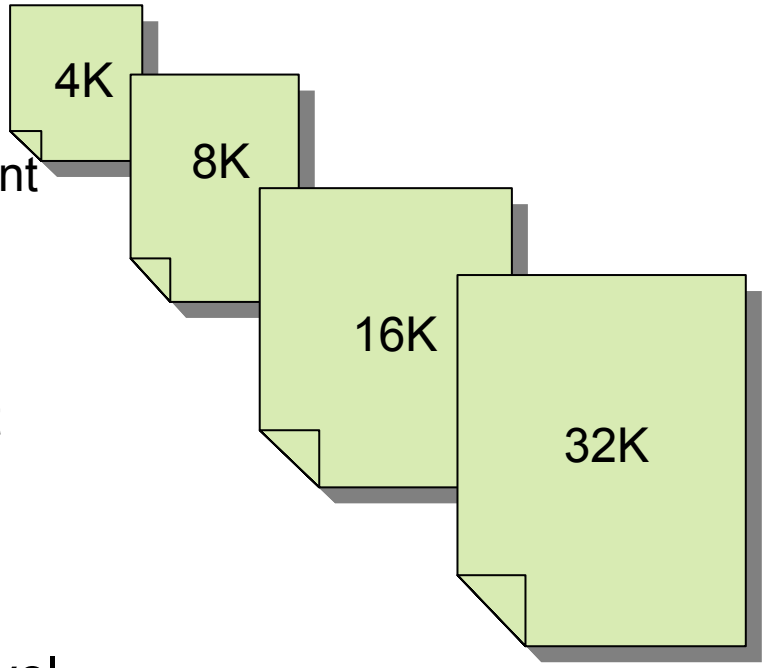

- Larger index pages allow for more efficient use of storage
  - Fewer page splits for long keys
  - More key values per page
- Define RANDOM index keys to avoid hot spots with multiple processes inserting sequential keys
- Rebuild Index SHRLEVEL CHANGE
- Index compression provides page-level compression
  - Data is compressed to 4K pages on disk
  - 32K/16K/8K pages results in up to 8x/4x/2x disk savings
  - No compression dictionaries
    - Compression on the fly
    - No LOAD or REORG required

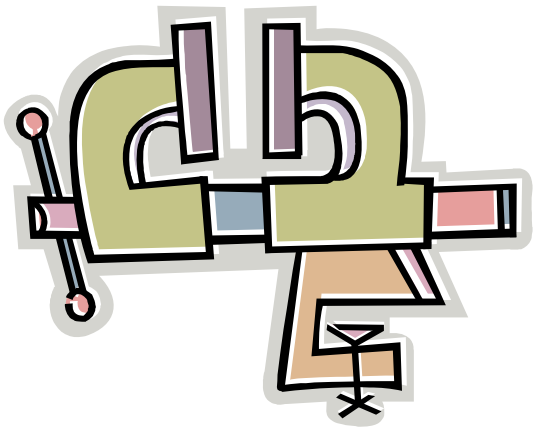

Asymmetric Index Page Splits

Multiple Sequential Insert Patterns on an Index

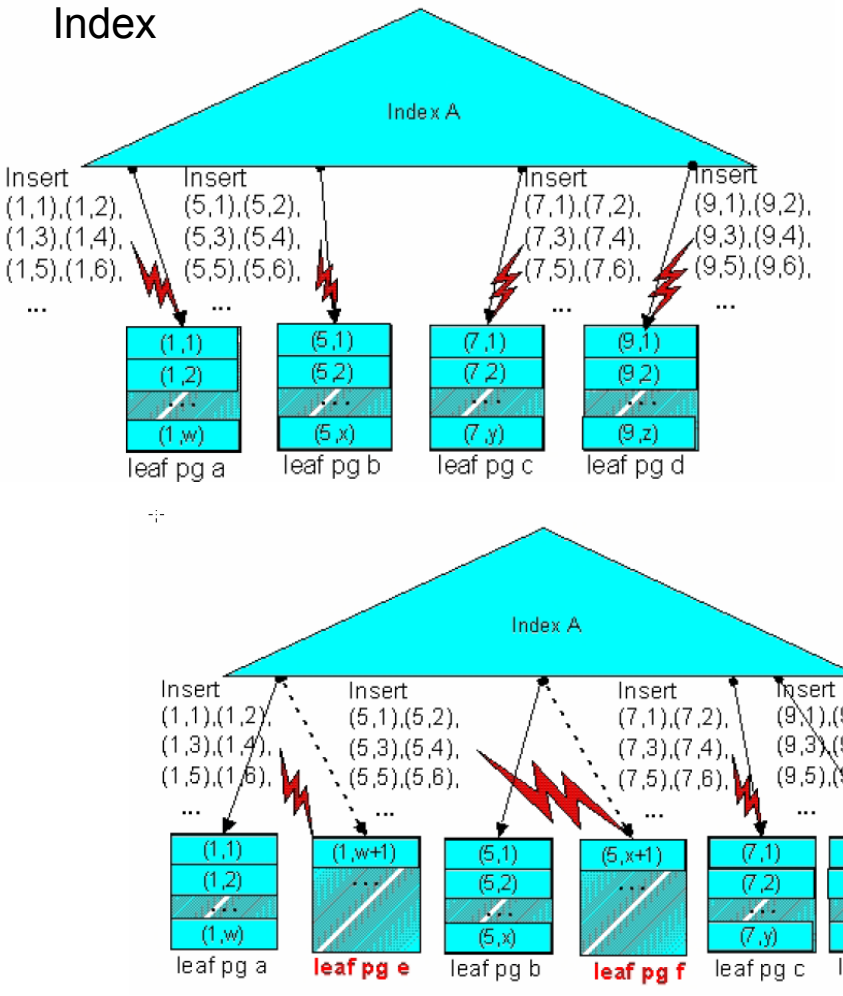

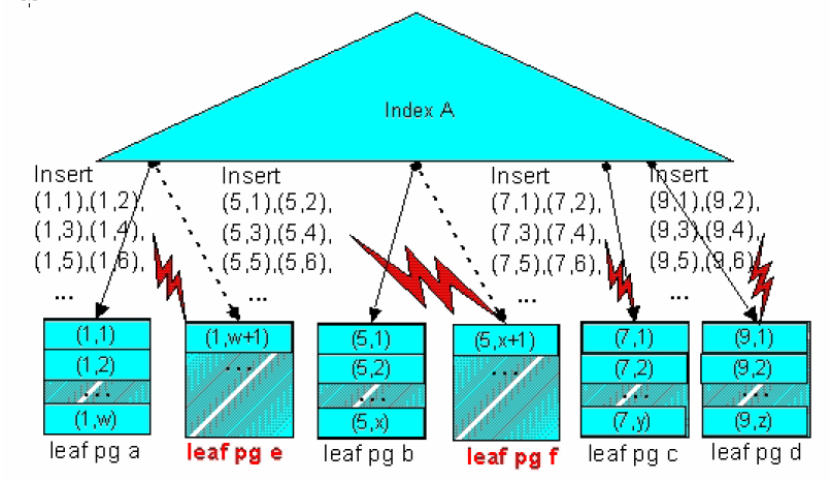

Sequential inserts into the middle of an index resulted in some pages with 50% free space prior to V9

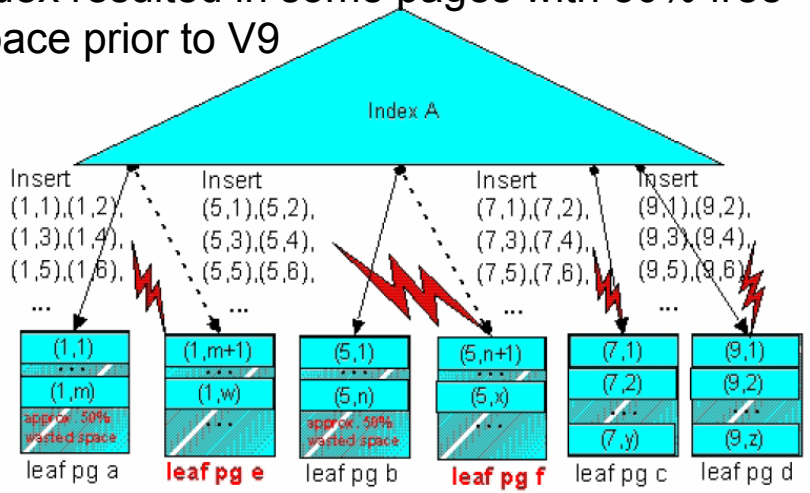

New algorithm dynamically accommodates a varying pattern of inserts

Index on Expression

- Simple indexes can contain concatenated columns

```
create index totalComp on
  W2_TABLE(salary, bonus)
```

- Index on expression
  - Value of the index has been transformed
  - May not be the value of any of the columns that it is derived from
  - Optimizer can use this index

```
Create index totalComp on
  W2_TABLE(salary+bonus)
```

W2_TABLE

name	salary	bonus
Gary	20,000	500
Paul	40,000	5,000
Matt	400,000	10,000
Rex	40,000	2,000
Rob	200,000	210,000

Query / Order on Total Compensation

Data Sharing

Data Sharing V9 Enhancements

- Log latch contention relief
- Restart performance enhancements
  - Reduced impact of retained locks – released as rollbacks are completed
  - Open data sets ahead of log apply
- Command to remove GBP-dependency at object level
  - ACCESS DB MODE(NGBPDEP)
  - Typical usage would be before batch run
  - Command to “prime” open data set
  - ACCESS DB MODE(OPEN) [PART]
- Auto-recover GRECP/LPL objects on group restart
  - Useful in Disaster Recovery or GDPS scenarios
- DB2 overall health taken into account for WLM routing
- Balance group attach connections across multiple members on same LPAR (V7, V8 usermod)

Utilities

Utilities Highlights

- More online utilities
  - Rebuild Index SHRLEVEL CHANGE
  - Reorg LOB now supports SHRLEVEL REFERENCE (space reclamation)
  - Check data, LOB and repair locate … SHRLEVEL CHANGE
  - Check index SHRLEVEL REFERENCE supports parallel for > 1 index
  - Clones for “online LOAD REPLACE”
- Online REORG BUILD2 phase elimination
- REORG parallelism for UNLOAD, RELOAD, LOG phases
- Utility TEMPLATE switching
- UNLOAD SKIP LOCKED DATA option

Utilities Highlights…

- MODIFY Recovery enhancements
  - "*Retain*" keyword added to improve management of copies
    - LAST(n), LOGLIMIT, GDGLIMIT
- Volume-based COPY/RECOVER (BACKUP SYSTEM/RESTORE SYSTEM)
  - RECOVER modified to enable object-level recovery from volume FlashCopy
  - Full integration of tape into BACKUP/RESTORE SYSTEM utilities
  - Incremental FlashCopy, APAR PK41001
- Truncate log based on timestamp
- RECOVER to any point-in-time with consistency
- RECOVER RESTOREBEFORE to use an earlier image copy
- Display progress of RECOVER during log apply
- COPY CHECKPAGE option always active
  - "Copy Pending" avoided if broken page encountered
- COPY SCOPE PENDING to copy only objects in "Copy Pending"

Utilities Performance Improvements

- CPU reductions in LOAD (with additional savings if data is PRESORTED), REORG, and REBUILD
  - Reductions mostly due to improved index processing (* with exceptions)
    - 10 to 20% in Image Copy* (even with forced CHECKPAGE YES)
    - 5 to 30% in Load, Reorg, Reorg Partition, Rebuild Index
      - Except REORG TABLESPACE SHR CHG PART with NPSIs
    - 20 to 40% in Load
    - 20 to 60% in Check Index
    - 35% in Load Partition
    - 30 to 40% in Runstats Index
    - 40 to 50% in Reorg Index
    - Up to 70% in Load Replace Partition with dummy input

And Also …

- RLF enhancements
- Many serviceability enhancements
  - Messages
  - Display commands
  - Refresh ERLY code without IPL
  - Many others …
- Converged TEMP Space
- Zparm to prevent workfile monopolization
- SQL
  - FETCH FIRST/ ORDER BY in Subselect
- RTS moved to catalog – automatic update
- Global Query Optimization will allow DB2 to optimize a query as a whole rather than as independent parts
  - considers the effect of one query block on another
  - considers reordering query blocks
- Serviceability
  - DISPLAY enhancements
  - Health Monitoring – internal monitor
- V9 can use more index lookaside for INSERT & DELETE
- SSL
- Improved trace filtering
- Command line processor

Planned V9 Post GA Deliveries

- Text search server
- New XMLTABLE and XLMCAST functions:
  - APAR PK51573
- Incremental FlashCopy
  - APAR PK41001, z/OS 1.8 APAR OA17314
- Plan stability
  - APAR PK52523
- Trusted context enhancements
  - APAR PK44617
- New storage class Zparm for online CHECK utilities
  - APAR PK41711. Needed for when PPRC is used.
- ALTER TABLE ALTER COLUMN DROP DEFAULT
- Allow RESTORE SYSTEM recovery without requiring log truncation
  - APAR PK51979
- Spatial data type enhancements

Futures: Possible Items in DB2 Vnext

DB2 for z/OS Into the Future

Delivering Customer Value

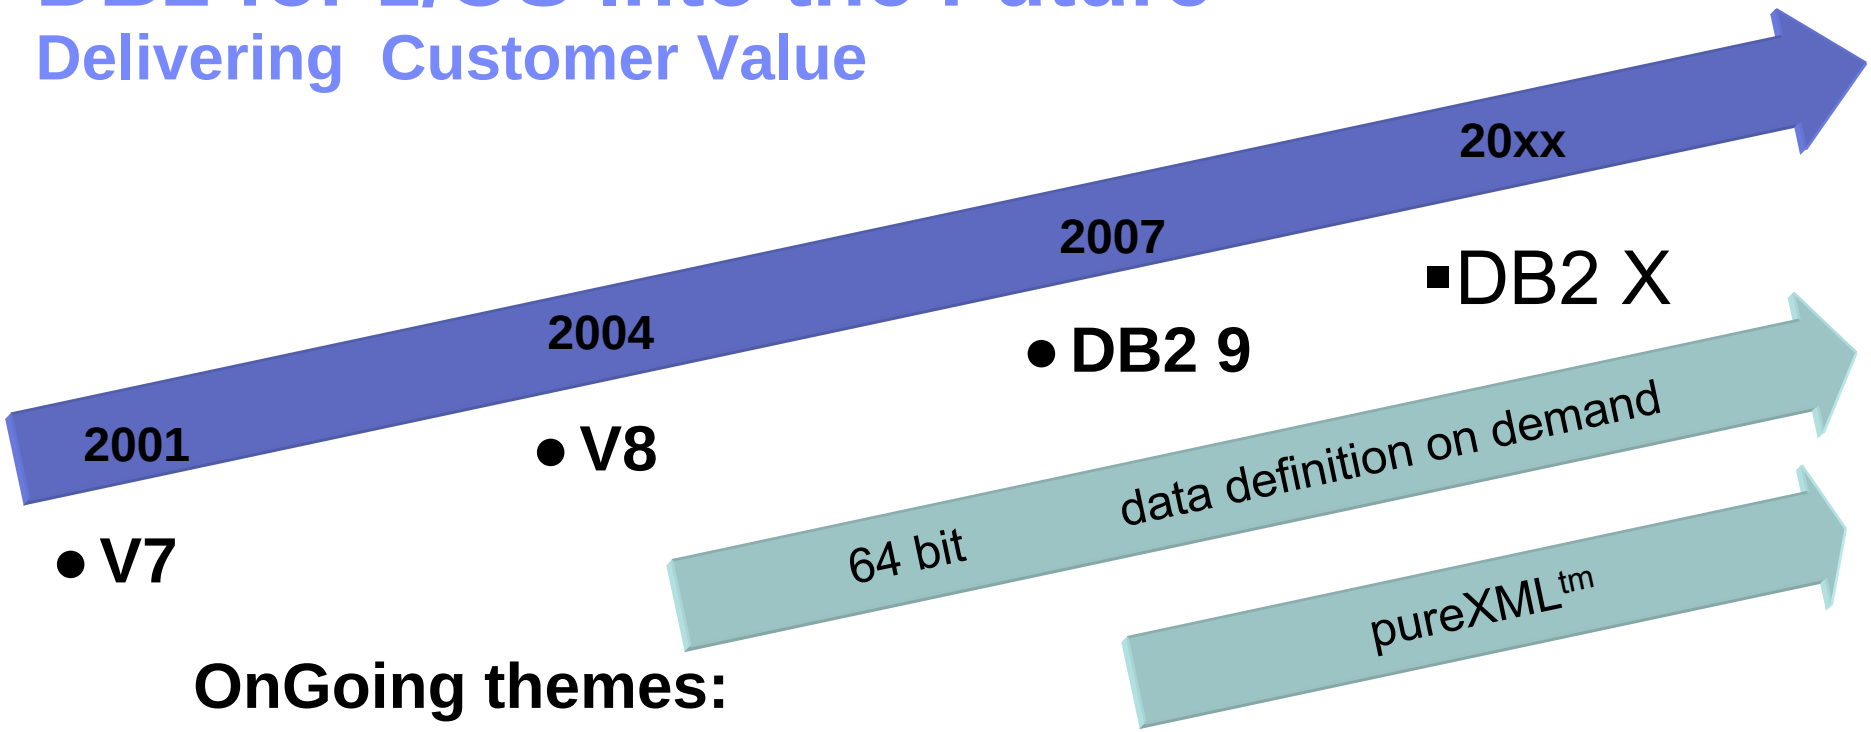

OnGoing themes:

Performance Scalability
Reliability Availability Serviceability
Security Productivity
Application Development
SQL XML SOA

DB2 X for z/OS Status

- The following slides represent DB2 Development's current thinking on some of the items that are candidates for Vnext
- It is still early in the development process, so the details will change
- The intention is to give you some information on DB2's future technical directions
- DB2 Development values customer feedback

64-bit Evolution – The Final Frontier

- DB2 vstor constraint is still an important issue for some customers
  - Can limit number of concurrent active threads on a single DB2
  - Can make it more difficult to grow DB2 environments
    - Consolidation through mergers / acquisitions
    - Enterprise workload consolidation to reduce costs
    - Some customers put multiple DB2 members from the same DSG on the same LPAR
  - Detailed monitoring of IFCIDs 225 and 217 sometimes necessary
- Vnext: dramatic reduction of DBM1 31-bit private vstor
  - Allow for big increase in max number of active threads
  - Remove need for detailed DBM1 virtual storage monitoring
  - May allow for consolidation to fewer DB2 members on fewer LPARs
  - Reduced cost, easier to manage, easier to grow

Other System Scaling Improvements

- Other bottlenecks can emerge in extremely heavy workloads
  - Vnext: several improvements planned to reduce latching and other system serialization contention
  - Vnext: new ISO(CS) option to avoid waiting on uncommitted updates
- Concurrent DDL/BIND/Prepare processes can hit contention with one another
  - Vnext: restructure parts of the DB2 catalog to avoid the contention

Performance, Performance, Performance

- DB2 applications require always-improving performance
- Future machines require new s/w performance techniques
  - Higher n-way SMPs, more memory, higher cache-miss penalties
- Vnext objective: significantly improved DB2 performance for a wide range of applications
  - New high performance technique for p-key access
  - Make DB2 code and control structures more cpu cache friendly
  - Buffer pool enhancements
  - More indexing enhancements
  - Inline LOBs, LOB streaming inside DB2
  - DDF performance enhancements
  - Several query performance improvement items

RAS: extend DB2's industry leadership

- Continuous availability requirements continue to escalate
- Nice big batch windows are a thing of the past
- DBAs increasingly need the ability to make all changes and to do all maintenance activities online
- Vnext:
  - Allow for more online schema changes: table space type, page size, some others
  - REORG SHRLEVEL(CHANGE) for LOBs
  - Consistent image copy without quiesce
  - Inline copies to allow for dataset-level FlashCopy
  - Online REORG usability and performance enhancements

Security

- Regulatory compliance initiatives and auditing requirements increase the need for data retention and tightening up data base security
  - Time-series queries are often difficult
  - Views don't tend to work well with lots of user sets all with different privileges
- Vnext:
  - New table attributes to allow for retention of older versions of data.
    - "snapshot" query capability to see what data looked like at a particular point in time
  - Fine grained access control
    - Rules to restrict end user access to individual cells
    - Allow specification of value returned with access denied, e.g. xxxx-xxxxxxx-1234 (credit card number)
    - Mandatory, cannot be bypassed like views
    - Avoid cost of maintaining views, transparent to applications
  - Finer granularity SYSADM and DBADM privileges

Productivity – Doing More with Less!

- IT staffs continue to be stretched thinner
- DB2 data bases and workloads continue growing larger
- It's often difficult to know when to run RUNSTATS or REORG or when to activate/deactivate compression
- Vnext:
  - Automatic stats collection
  - Automatic enable/disable compression
  - Safe query optimization and plan stability
  - REORG avoidance: avoid impact of disorganized indexes
  - Autonomic diagnosis and tuning for query performance issues
  - Automatic configuration of IBM supplied SPs and UDFs
  - DROP PARTITION, ALTER ADD PARTITION enhancements

Application Development

- Apps need the ability to embed logic in triggers and to use SQLPL in user defined functions
- Apps need to specify BIFs or UDFs in check constraints
- V9 introduced DECFLOAT but it can't be indexed
- Dynamic SQL statements with literals can't be cached
- Vnext:
  - SQLPL in triggers and UDFs (scalar and table functions)
  - Allow BIFs or scalar UDFs in check constraint definition
  - Indexing support for DECFLOAT + performance improvements
  - Allow caching of dynamic SQL statements with literals
  - Instance based statement hints to avoid app changes for opt hints

SQL

- Porting from other platforms to DB2 for z/OS is sometimes difficult due to SQL differences
- Apps often need new, enhanced SQL capabilities
- Vnext:
  - Session variables to provide a flexible way to pass data from one program to the next via SQL – CREATE VARIABLE
  - Generated column support
    - ALTER TABLE employee
      ADD COLUMN upper_lastname
      GENERATED ALWAYS AS
      (UPPER(lastname));
  - Greater TIMESTAMP precision
  - Special indicator for “value not supplied” or “DEFAULT”
  - SQL Paging for when only a specific part of the result set is needed

Data Warehousing

- DB2 for z/OS is often attractive for warehousing, but further improvements are required for some apps
- New SQL functions are required for more advanced analytics
- DB2 restricts use of query parallelism for certain cases
- Vnext:
  - Moving Sum, Moving Average
  - N-tile functions
  - Enhanced query parallelism technology for improved performance
    - Remove query parallelism restrictions
  - Dynamic Index ANDing improvements
  - In-memory techniques for faster query performance

Service Oriented Architecture and XML

- Java environments require simplified monitoring and problem determination for end-to-end data base access
- Stored procs are often exposed as web services, but certain function is lacking in some cases
- Vnext:
  - Support for pureQuery extensions to Java
  - Prepare attribute that adds optimistic locking to SELECT list
  - RETURN TO CLIENT capability for nested stored procedures

Service Oriented Architecture and XML…

- Native XML datatype was introduced in V9 and some enhancements are required for some applications
- Vnext:
  - XML schema validation in the engine for improved usability, performance
  - Binary XML exchange format for improved performance
  - XML multi-versioning for more robust XML queries
  - Allow easy update of sub-parts of an XML document
  - Introduction of XQuery syntax
  - Stored proc, UDF, Trigger enhanced support for XML

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

IBM Information On Demand 2007

IBM

Act.Right.Now.

IBM INFORMATION ON DEMAND 2007
October 14 – 19, 2007
Mandalay Bay
Las Vegas, Nevada

Tooling

Tooling

- Developer Workbench (DWB)
- DB2 9 .Net Add-Ins for Visual Studio 2005
- Optimization Service Center
- Optimization Expert
- Visual Explain – to be integrated with IBM Data Studio

Developer Workbench

- Create and manage database objects
  - Create tables, indexes, view, MQTs, UDFs…
  - Create SQL
  - Browse, extract and load tables
- Stored Procedures
  - Create, edit, debug, deploy and testing DB2 stored procedures
- Deploy routines to unlike servers
- Manage and share project resources
- XML Features:
  - Edit XML documents
  - Edit XML table data
  - Extract / Load XML data

Optimization Service Center and Optimization Expert

- DB2 for z/OS Optimization Service Center
- DB2 Optimization Expert for z/OS V1.1

Functions	Visual Explain	Optimization Service Center	Optimization Expert
Queries from Cache, Catalog	Yes	Yes	Yes
Query Formatter, Annotation		Yes	Yes
Access Plan Graph	Yes	Yes	Yes
Visual Plan Hint		Yes	Yes
Query Statistics Advisor	Yes	Yes	Yes
Workload Statistics Advisor		Yes	Yes
Profile based Monitoring		Yes	Yes
Query Index Advisor			Yes
Workload Index Advisor			Yes
Query Advisor			Yes
Workload Query Advisor			Yes
Access Path Advisor			Yes